NOKIA

7450 ETHERNET SERVICE SWITCH
7750 SERVICE ROUTER
VIRTUALIZED SERVICE ROUTER

TRIPLE PLAY SERVICE DELIVERY ARCHITECTURE GUIDE RELEASE 16.0.R4

3HE 14143 AAAB TQZZA 01

Issue: 01

November 2018

Nokia — Proprietary and confidential.
Use pursuant to applicable agreements.

Nokia is a registered trademark of Nokia Corporation. Other products and company names mentioned herein may be trademarks or tradenames of their respective owners.

The information presented is subject to change without notice. No responsibility is assumed for inaccuracies contained herein.

© 2018 Nokia.

Contains proprietary/trade secret information which is the property of Nokia and must not be made available to, copied or used by anyone outside Nokia without its written authorization. Not to be used or disclosed except in accordance with applicable agreements.

Table of Contents

1 Getting Started 29
1.1 About This Guide 29
1.2 Nokia SR OS Services Configuration Process 31

2 Introduction to Triple Play 33
2.1 Nokia's Triple Play Service Delivery Architecture 33
2.1.1 Introduction to Triple Play 33
2.1.2 Blueprint for Optimizing Triple Play Service Infrastructures 34
2.1.3 Architectural Foundations 35
2.1.4 Optimizing Triple Play Service Infrastructures 36
2.1.4.1 Distributed Service Edges 38
2.1.4.2 Service Differentiation, QoS Enablement 40
2.1.4.3 Virtual MAC Subnetting for VPLS 43
2.2 Services 47
2.2.1 Service Types 47
2.2.2 Service Policies 48
2.3 Nokia Service Model 50
2.3.1 Introduction 50
2.3.2 Service Entities 50
2.3.3 Customers 51
2.3.4 Service Access Points (SAPs) 51
2.3.4.1 SAP Encapsulation Types and Identifiers 52
2.3.4.2 Ethernet Encapsulations 52
2.3.4.3 SAP Considerations 53
2.3.5 Service Distribution Points (SDPs) 54
2.3.5.1 SDP Binding 55
2.3.5.2 Spoke and Mesh SDPs 55
2.3.5.3 SDP Encapsulation Types 56
2.3.5.4 SDP Keepalives 57
2.4 Epipe Service Overview 59
2.5 VPLS Service Overview 60
2.5.1 Split Horizon SAP Groups and Split Horizon Spoke SDP Groups 61
2.5.1.1 Residential Split Horizon Groups 61
2.6 IES Service Overview 62
2.6.1 IP Interface 62
2.7 VPRN Service Overview 64
2.8 Deploying and Provisioning Services 65
2.8.1 Phase 1: Core Network Construction 65
2.8.2 Phase 2: Service Administration 65
2.8.3 Phase 3: Service Provisioning 65
2.9 Configuration Notes 66
2.9.1 General 66
2.10 Configuring Triple Play Services with CLI 67
2.10.1 Configuring VPLS Residential Split Horizon Groups 67
2.10.1.1 Configuring Static Hosts 67

2.11 Triple Play Services Command Reference 73
2.11.1 Command Hierarchies....73
2.11.1.1 IES/VPRN Triple Play Commands 73
2.11.1.2 VPLS Triple Play Commands....77
2.11.1.3 Show Commands 82
2.11.1.4 Clear Commands....84
2.11.1.5 Debug Commands....85
2.11.1.6 Tools Commands 85
2.11.2 Command Descriptions 85
2.11.2.1 Global Commands....85
2.11.2.2 IES/VPRN Service Commands 87
2.11.2.3 Show Commands 162
2.11.2.4 Clear Commands....267
2.11.2.5 Debug Commands....272
2.11.2.6 Tools Commands 274

3 DHCP Management....275
3.1 DHCP Principles....275
3.2 DHCP Features 278
3.2.1 DHCP Relay 278
3.2.2 DHCPv4 Relay Proxy 278
3.2.3 Subscriber Identification Using Option 82 Field....286
3.2.3.1 Trusted and Untrusted....287
3.2.4 DHCP Snooping 288
3.2.5 DHCP Lease State Table 289
3.2.6 DHCP and Layer 3 Aggregation 291
3.2.6.1 DHCPv4 Snooping 291
3.2.6.2 DHCPv6 Snooping 293
3.2.7 Local DHCP Servers 295
3.2.7.1 Overview....295
3.2.7.2 Local DHCP Server Support....297
3.2.8 DHCPv6....298
3.2.8.1 DHCPv6 Relay Agent....298
3.2.8.2 DHCPv6 Prefix Options 299
3.2.8.3 Neighbor Resolution With DHCPv6 Relay....299
3.2.8.4 DHCPv6 Lease Persistency 299
3.2.8.5 Local Proxy Neighbor Discovery 299
3.2.8.6 IPv6oE Hosts Behind Bridged CPEs 300
3.2.8.7 IPv6 Link-Address Based Pool Selection 300
3.2.8.8 IPv6 Address and Prefix Stickiness....300
3.2.8.9 IPv4/v6 Linkage for Dual-Stack Hosts or Layer 3 RGs....301
3.2.8.10 Host Connectivity Checks for IPv6 301
3.2.9 Lease Query....301
3.2.10 DHCPv6 to Server Option 302
3.2.11 Flexible Host Identification in LUDB Based on DHCPv4/v6 Options 302
3.2.12 DHCP Caching 303
3.2.13 Flexible Creation of DHCPv4/6 Host Parameters 304
3.2.14 Python DTC Variables and API 304

3.2.14.1 DTC Debugging Facility....309
3.2.15 Virtual Subnet for DHCPv4 Hosts....310
3.2.16 Address Reservation for Sticky Leases....310
3.3 Proxy DHCP Server....312
3.3.1 Local DHCP Servers....317
3.3.1.1 Terminology....317
3.3.1.2 Overview....318
3.3.1.3 DHCP Lease Synchronization....322
3.3.1.4 Intercommunication Link Failure Detection....323
3.3.1.5 DHCP Server Failover States....324
3.3.1.6 Lease Time Synchronization....325
3.3.1.7 Maximum Client Lead Time (MCLT)....327
3.3.1.8 Sharing IPv4 Address-Range or IPv6 Prefixes....330
3.3.1.9 Fast-Switchover of IP Address and Prefix Delegation....334
3.4 Local Address Assignment....336
3.4.1 Stateless Address Auto-Configuration....336
3.5 Configuring DHCP with CLI....337
3.5.1 Enabling DHCP Snooping....337
3.5.2 Configuring Local User Database Parameters....338
3.5.3 Configuring Option 82 Handling....346
3.5.4 Enabling DHCP Relay....346
3.6 Triple Play DHCP Command Reference....349
3.6.1 Command Hierarchies....349
3.6.1.1 System DHCP Commands....349
3.6.1.2 DHCP Filter Commands....349
3.6.1.3 Router DHCP Local User Database Commands....350
3.6.1.4 Local User Database Commands....353
3.6.1.5 Service DHCP Commands....360
3.6.1.6 Show Commands....372
3.6.1.7 Tools Commands....373
3.6.1.8 Clear Commands....374
3.6.1.9 Debug Commands....375
3.6.2 Command Descriptions....376
3.6.2.1 Triple Play DHCP Configuration Commands....376

4 Stateless Address Auto-configuration (SLAAC)....573
4.1 SLAAC Management Principles....573
4.2 Configuration Overview....573
4.3 Router-Solicit Trigger....574
4.4 SLAAC Address Assignment....574
4.5 Static SLAAC Prefix Assignment....574
4.6 Dynamic SLAAC Prefix Assignment....575
4.7 SLAAC Prefix Replacement....576

5 Point-to-Point Protocol over Ethernet Management....577
5.1 PPPoE....577
5.1.1 PPPoE Authentication and Authorization....580
5.1.1.1 General Flow....580
5.1.1.2 RADIUS....581

5.1.1.3 Local User Database Directly Assigned to PPPoE Node 581
5.1.1.4 Subscriber per PPPoE Session Index 582
5.1.1.5 Local DHCP Server with Local User Database 585
5.1.2 Multiple Sessions Per MAC Address 586
5.1.3 PPP Session Re-establishment 587
5.1.4 Private Retail Subnets 587
5.1.5 IPCP Subnet Negotiation 589
5.1.5.1 Numbered WAN Support for Layer 3 RGs 590
5.1.6 IES as Retail Service for PPPoE Host 590
5.1.7 Unnumbered PPPoX 591
5.1.8 Selective Backhaul of PPPoE Traffic using an Epipe Service 592
5.2 MLPPPoE, MLPPP(oE)oA with LFI on LNS 595
5.2.1 Terminology 595
5.2.2 LNS MLPPPoX 595
5.2.3 MLPPP Encapsulation 596
5.2.4 MLPPPoX Negotiation 597
5.2.5 Enabling MLPPPoX 598
5.2.6 Link Fragmentation and Interleaving (LFI) 599
5.2.6.1 MLPPPoX Fragmentation, MRRU and MRU Considerations 600
5.2.7 LFI Functionality Implemented in LNS 602
5.2.7.1 Last Mile QoS Awareness in the LNS 604
5.2.7.2 BB-ISA Processing 606
5.2.7.3 LNS-LAC Link 607
5.2.7.4 AN-RG Link 608
5.2.7.5 Home Link 608
5.2.7.6 Optimum Fragment Size Calculation by LNS 608
5.2.8 Upstream Traffic Considerations 611
5.2.9 Multiple Links MLPPPoX With No Interleaving 611
5.2.10 MLPPPoX Session Support 611
5.2.11 Session Load Balancing Across Multiple BB-ISAs 614
5.2.12 BB-ISA Hashing Considerations 614
5.2.13 Last Mile Rate and Encapsulation Parameters 615
5.2.14 Link Failure Detection 618
5.2.15 CoA Support 618
5.2.16 Accounting 619
5.2.17 Filters and Mirroring 620
5.2.18 PTA Considerations 620
5.2.19 QoS Considerations 620
5.2.19.1 Dual-Pass 620
5.2.19.2 Traffic Prioritization in LFI 621
5.2.19.3 Shaping Based on the Last Mile Wire Rates 622
5.2.19.4 Downstream Bandwidth Management on Egress Port 624
5.2.20 Sub/Sla-Profile Considerations 624
5.2.21 Example of MLPPPoX Session Setup Flow 625
5.2.22 Other Considerations 627
5.3 Configuration Notes 628
5.4 PPP Command Reference 631
5.4.1 Command Hierarchies 631
5.4.1.1 PPPoE Policy Configuration Commands 631

5.4.1.2 PPPoE Service Commands....632
5.4.1.3 PPPoE Local User Database Commands....633
5.4.1.4 MLPPP on LNS Commands....634
5.4.1.5 Show Commands....637
5.4.2 Command Descriptions....637
5.4.2.1 Generic Commands....637
5.4.2.2 Service Commands....669
5.4.2.3 Show Commands....670
5.4.2.4 Clear Commands....681
5.4.2.5 Debug Commands....682

6 Layer 2 Tunneling Protocol (L2TP)....687
6.1 Terminology....687
6.2 L2TP Overview....687
6.2.1 LAC DF Bit....687
6.2.2 Handling L2TP Tunnel/Session Initialization Failures....688
6.2.2.1 L2TP Tunnel/Session Initialization Failover Mechanisms on LAC....688
6.2.2.2 Peer Blacklist....689
6.2.2.3 Tunnel Blacklists....689
6.2.2.4 Tunnel Selection Mechanism....692
6.2.2.5 Tunnel Probing....692
6.2.2.6 Controlling the Size of Blacklist....693
6.2.2.7 Displaying the Content of a Blacklist....693
6.2.2.8 Generating Trap when the Blacklist is Full....694
6.2.2.9 Premature Removal of Blacklisted Entries....694
6.2.2.10 Manual Purging of Entities within the Blacklist....694
6.2.3 CDN Result Code Overwrite....695
6.2.4 LNS Proxy....695
6.3 L2TP LAC VPRN....697
6.3.1 Per-ISP Egress L2TP DSCP Reclassification....699
6.4 Traffic Steering on L2TP LAC....701
6.4.1 Steering Activation and Deactivation....702
6.4.2 Steering States....702
6.4.3 Configuring Traffic Steering on L2TP LAC....703
6.5 L2TP Tunnel RADIUS Accounting....705
6.5.1 Accounting Packets List....706
6.5.2 RADIUS Attributes Value Considerations....709
6.6 MLPPP on the LNS Side....710
6.7 LNS Reassembly....710
6.8 LNS Subscriber Policers....712
6.8.1 Policer Support....712
6.9 L2TP Command Reference....713
6.9.1 Command Hierarchies....713
6.9.1.1 L2TP Configuration Commands....713
6.9.1.2 Steering Profile Commands....715
6.9.1.3 L2TP Tunnel RADIUS Accounting Commands....716
6.9.1.4 Show Commands....716
6.9.1.5 Clear Commands....717
6.9.1.6 Debug Commands....717

6.9.1.7 Tools Commands 718
6.9.2 Command Descriptions 719
6.9.2.1 L2TP Configuration Commands 719
6.9.2.2 Steering Profile Commands 733
6.9.2.3 L2TP Tunnel RADIUS Accounting Commands 735
6.9.2.4 Show Commands 741
6.9.2.5 Debug Commands 749

7 Triple Play Security 753
7.1 Triple Play Security Features 753
7.1.1 Anti-Spoofing Filters 753
7.1.1.1 Anti-spoofing Filter Types 754
7.1.1.2 Filtering Packets 754
7.1.2 Layer 2 Triple Play Security Features 755
7.1.2.1 MAC Pinning 755
7.1.2.2 MAC Protection 755
7.1.2.3 DoS Protection 756
7.1.2.4 VPLS Redirect Policy 758
7.1.3 ARP Handling 759
7.1.3.1 ARP Reply Agent 759
7.1.3.2 Dynamic ARP Table Population 760
7.1.3.3 Local Proxy ARP 760
7.1.4 Web Portal Redirect 761
7.2 Configuring Triple Play Security with CLI 765
7.2.1 Common Configuration Tasks 765
7.2.1.1 Configuring Anti-Spoofing Filters 765
7.2.1.2 Configuring Triple Play Security Features 765
7.2.1.3 Configuring ARP Handling 769
7.2.1.4 Configuring Web Portal Redirect 771
7.3 Triple Play Security Command Reference 775
7.3.1 Command Hierarchies 775
7.3.1.1 Anti-Spoofing Commands 775
7.3.1.2 Layer 2 Security Commands 775
7.3.1.3 ARP Handling Commands 776
7.3.1.4 Show Commands 776
7.3.1.5 Tools Commands 777
7.3.2 Command Descriptions 777
7.3.2.1 Triple Play Security Configuration Commands 777
7.3.2.2 Show Commands 788
7.3.2.3 Tools Commands 789

8 Triple Play Multicast 793
8.1 Introduction to Multicast 793
8.2 Multicast in the Broadband Service Router 794
8.2.1 Internet Group Management Protocol 794
8.2.1.1 IGMP Versions and Interoperability Requirements 794
8.2.1.2 IGMP Version Transition 795
8.2.2 Multicast Listener Discovery 796
8.2.2.1 MLD Versions and Interoperability Requirements 796

8.2.2.2 Source Specific Multicast Groups....797
8.2.3 Protocol Independent Multicast Sparse Mode (PIM-SM)....797
8.2.4 Ingress Multicast Path Management (IMPM) Enhancements....799
8.3 Multicast in the BSA....801
8.3.1 IGMP Snooping....801
8.3.1.1 IGMP/MLD Message Processing....802
8.3.1.2 IGMP Message Processing....803
8.3.1.3 MLD Message Processing....804
8.3.1.4 IGMP/MLD Filtering....804
8.3.2 Multicast VPLS Registration (MVR)....804
8.3.3 Layer 3 Multicast Load Balancing....805
8.3.4 IGMP State Reporter....807
8.3.4.1 IGMP Data Records....808
8.3.4.2 Transport Mechanism....811
8.3.4.3 HA Compliance....811
8.3.4.4 QoS Awareness....811
8.3.4.5 IGMP Reporting Caveats....811
8.4 Multicast Support over Subscriber Interfaces in Routed CO Model....812
8.4.1 Multicast Over IPoE....814
8.4.1.1 Per SAP Replication Mode....815
8.4.1.2 Per Subscriber Host Replication Mode....824
8.4.2 Multicast Over PPPoE....828
8.4.3 IGMP Flooding Containment....829
8.4.4 IGMP/MLD Timers....830
8.4.5 IGMP/MLD Query Intervals....830
8.4.6 HQoS Adjustment....831
8.4.6.1 Host Tracking (HT) Considerations....836
8.4.6.2 HQoS Adjust Per Vport....837
8.4.7 Redirection....841
8.4.8 Hierarchical Multicast CAC (H-MCAC)....842
8.4.8.1 MCAC Bundle Bandwidth Limit Considerations....845
8.4.9 Determining MCAC Policy in Effect....849
8.4.10 Multicast Filtering....850
8.4.11 Joining the Multicast Tree....852
8.4.12 Wholesale/Retail Requirements....852
8.4.13 QoS Considerations....853
8.4.14 Redundancy Considerations....854
8.4.14.1 Redirection Considerations....857
8.4.15 Query Intervals for Multicast....857
8.4.15.1 ESM Host-based Queries....858
8.4.15.2 Group Interface-based Queries....859
8.5 Configuring Triple Play Multicast Services with CLI....861
8.5.1 Configuring IGMP Snooping in the BSA....861
8.5.1.1 Enabling IGMP Snooping in a VPLS Service....861
8.5.1.2 IGMPv3 Multicast Routers....861
8.5.1.3 With IGMPv1/2 Multicast Routers....861
8.5.1.4 Modifying IGMP Snooping Parameters....862
8.5.1.5 Modifying IGMP Snooping Parameters for a SAP or SDP....862
8.5.2 Configuring Static Multicast Groups on a SAP or SDP....864

8.5.2.1 Enabling IGMP Group Membership Report Filtering 865
8.5.3 Configuring Multicast VPLS Registration (MVR) 867
8.5.4 Configuring IGMP, MLD, and PIM in the BSR 868
8.5.4.1 Enabling IGMP 868
8.5.4.2 Configuring IGMP Interface Parameters 869
8.5.4.3 Configuring Static Parameters 869
8.5.4.4 Configuring SSM Translation 870
8.5.4.5 Enabling MLD 871
8.5.4.6 Configuring MLD Interface Parameters 871
8.5.4.7 Configuring Static Parameters 871
8.5.4.8 Configuring SSM Translation 872
8.5.4.9 Configuring PIM 873
8.5.4.10 Configuring Bootstrap Message Import and Export Policies 879
8.6 Triple Play Multicast Command Reference 881
8.6.1 Command Hierarchies 881
8.6.1.1 Multicast Management Commands 881
8.6.1.2 Multicast Info Policy Bundle Commands 882
8.6.1.3 Multicast Service Commands 883
8.6.1.4 Ingress Multicast Path Management Commands 886
8.6.1.5 Multicast Redirection 887
8.6.1.6 Multicast Query Timer Commands 888
8.6.1.7 Show Commands 889
8.6.1.8 Debug Commands 889
8.6.2 Command Descriptions 889
8.6.2.1 Multicast Management Configuration Commands 889
8.6.2.2 Show Commands 938
8.6.2.3 Debug Commands 941

9 Triple Play Enhanced Subscriber Management 943
9.1 Uniform RADIUS Server Configuration 943
9.1.1 RADIUS Server Configuration 943
9.1.1.1 Uniform RADIUS Server Configuration (Preferred) 943
9.1.1.2 Legacy RADIUS Server Configuration 946
9.2 RADIUS Authentication of Subscriber Sessions 948
9.2.1 RADIUS Authentication Extensions 949
9.2.1.1 Triple Play Network with RADIUS Authentication 950
9.2.2 RADIUS Authorization Extensions 951
9.2.2.1 Calling-Station-ID 952
9.2.2.2 Subscriber Session Timeout 952
9.2.2.3 RADIUS Reply Message for PPPoE PAP/CHAP 955
9.2.2.4 SHCV Policy 955
9.2.3 radius-server-policy Retry Attempt Overview 955
9.2.4 AAA RADIUS Server Operation Status 956
9.2.5 AAA RADIUS Accounting Server Stickiness 958
9.2.6 AAA RADIUS Authentication Fallback Action 958
9.2.7 AAA Test User Account 959
9.2.8 Troubleshooting the RADIUS Server 960
9.2.9 Provisioning of Enhanced Subscriber Management (ESM) Objects 960
9.2.9.1 Provisioning IP Configuration of the Host 962

9.2.9.2 RADIUS-Based Authentication in Wholesale Environment 962
9.2.9.3 Change of Authorization and Disconnect-Request....963
9.2.9.4 RADIUS-Based Accounting....968
9.2.9.5 RADIUS Accounting Terminating Cause....973
9.2.9.6 Accounting Modes Of Operation 973
9.2.9.7 Per Session Accounting 976
9.2.9.8 RADIUS Session Accounting with PD as a Managed Route....977
9.2.9.9 No Host-Accounting....980
9.2.9.10 Host-Accounting Enabled 981
9.2.9.11 Reduction of Host Updates for Session Accounting Start and Stop....982
9.2.9.12 Accounting Interim Update Message Interval....982
9.2.9.13 CoA Triggered Accounting Interim Update....983
9.2.9.14 Class Attribute 983
9.2.9.15 User Name 984
9.2.9.16 Accounting-On and Accounting Off 984
9.2.9.17 RADIUS Accounting Message Buffering 988
9.2.9.18 Multiple Accounting Policies....991
9.2.9.19 Sending an Accounting Stop Message upon a RADIUS Authentication Failure of a PPPoE Session 991
9.3 Enhanced Subscriber Management Overview 993
9.3.1 Enhanced Subscriber Management Basics....993
9.3.1.1 Standard and Enhanced Subscriber Management....993
9.3.2 ESM for IPv6 1000
9.3.2.1 Models....1000
9.3.2.2 Setup 1002
9.3.2.3 64-bit and 128-bit WAN Mode 1003
9.3.2.4 Behavior 1007
9.3.2.5 Delegated-Prefix-Length....1012
9.3.2.6 DHCPv6 Relay Agent....1014
9.3.2.7 DHCPv6 Relay to Third Party DHCPv6 External Server....1016
9.3.2.8 DHCPv6 Local Server 1017
9.3.3 Dynamic Subscriber Host Processing 1018
9.3.3.1 Dynamic Tables....1018
9.3.4 Enhanced Subscriber Management Entities 1020
9.3.4.1 Instantiating a New Host....1021
9.3.4.2 Packet Processing for an Existing Host....1022
9.3.5 ESM Host Lockout 1023
9.3.5.1 Functionality 1024
9.3.6 ANCP and GSMP....1024
9.3.6.1 ANCP....1024
9.3.6.2 General Switch Management Protocol Version 3 (GSMPv3)....1026
9.3.6.3 DHCP Release Messages....1027
9.3.6.4 DHCP Client Mobility....1028
9.3.6.5 DHCP Lease Control....1028
9.3.7 Using Scripts for Dynamic Recognition of Subscribers 1028
9.3.7.1 Python Language and Programmable Subscriber Configuration Policy (PSCP)....1029
9.3.7.2 Determining the Subscriber Profile and SLA Profile of a Host 1029
9.3.7.3 Determining the Subscriber Profile....1031

9.3.7.4 Determining the SLA Profile 1032
9.3.8 Auto-Sub ID 1047
9.3.8.1 Sub-id Identifiers 1050
9.3.8.2 Dual Stack Hosts 1051
9.3.8.3 Mixing Hosts with Auto-Generated IDs and non Auto-Generated IDs 1051
9.3.8.4 PPPoA/PPPoEoA Considerations 1052
9.3.8.5 Deployment Considerations 1052
9.3.8.6 Caveats 1054
9.3.9 Limiting Subscribers and Hosts on a SAP 1054
9.3.10 Static Subscriber Hosts 1054
9.3.11 QoS for Subscribers and Hosts 1055
9.3.11.1 QoS Parameters in Different Profiles 1055
9.3.11.2 QoS Policy Overrides 1056
9.3.12 ESM Subscriber Hierarchical Traffic Control 1056
9.3.12.1 Subscriber HQoS 1056
9.3.12.2 Subscriber CFHP 1060
9.3.12.3 ATM/Ethernet Last-Mile Aware QoS for Broadband Network Gateway 1062
9.3.13 Subscriber Volume Statistics 1083
9.3.13.1 IP (Layer 3) Volume Accounting 1083
9.3.13.2 Separate IPv4 and IPv6 Counters 1084
9.3.14 Configuring IP and IPv6 Filter Policies for Subscriber Hosts 1090
9.3.14.1 IP Filter Attribute Format Details 1097
9.3.14.2 Checking Filter Policy Details 1102
9.3.15 ESM PPPoA/PPPoEoA 1103
9.3.15.1 PPPoA 1103
9.3.15.2 PPPoEoA 1105
9.3.15.3 Hardware Support 1106
9.3.15.4 Termination Points 1107
9.3.15.5 PPPoA Encapsulation 1107
9.3.15.6 Encapsulation Summary 1112
9.3.15.7 Concurrent Support for Different Service Types on the Same Port 1113
9.3.15.8 Restrictions in Scaled ATM MDA Mode 1113
9.3.15.9 QoS Implementation 1114
9.3.15.10 Association Between the Subscriber and ATM VC Traffic Descriptor (QoS) 1116
9.3.15.11 Per VP Shaping 1119
9.3.15.12 ATM/IOM QoS Integration 1121
9.3.15.13 Subscriber Instantiation Use Cases 1130
9.3.15.14 Authentication 1134
9.3.15.15 LUDB Access via Capture SAP 1135
9.3.15.16 Encapsulation Autosensing 1135
9.3.15.17 SAP Autoprovisioning 1136
9.3.15.18 PPP Nodes and ppp-policy 1137
9.3.15.19 MTU Considerations 1138
9.3.16 Multi-Chassis Synchronization 1140
9.3.16.1 Overview 1141
9.3.17 Subscriber Routed Redundancy Protocol (SRRP) 1143

9.3.17.1 SRRP Messaging 1143
9.3.17.2 SRRP and Multi-Chassis Synchronization 1144
9.3.17.3 SRRP Instance 1145
9.3.17.4 Subscriber Subnet Owned IP Address Connectivity 1148
9.3.17.5 Subscriber Subnet SRRP Gateway IP Address Connectivity 1148
9.3.17.6 Receive SRRP Advertisement SAP and Anti-Spoof 1149
9.3.18 PPPoE MC Redundancy 1149
9.3.18.1 SRRP Considerations for PPPoE 1149
9.3.18.2 State Synchronization 1151
9.3.18.3 Traffic Control and Redundant Interface 1153
9.3.18.4 MSAP Considerations 1157
9.3.18.5 Unnumbered Interface Support 1158
9.3.18.6 Compatibility with MC-LAG 1159
9.3.18.7 IPv6 Support 1159
9.3.18.8 Considerations with Local DHCP Server 1161
9.3.18.9 Redundant Interface Considerations 1162
9.3.19 Routed Central Office (CO) 1162
9.3.19.1 Layer 3 Subscriber Interfaces 1162
9.3.19.2 Wholesale Retail Routed CO 1169
9.3.19.3 Routed Subscriber Hosts 1175
9.3.20 Dual Homing 1185
9.3.20.1 Dual Homing to Two PEs (Redundant-Pair Nodes) in Triple Play Aggregation 1185
9.3.20.2 Steady-State Operation of Dual-homed Ring 1188
9.3.20.3 Broken-Ring Operation and the Transition to this State 1190
9.3.20.4 Transition from Broken to Closed Ring State 1192
9.3.20.5 Provisioning Aspects and Error Cases 1192
9.3.20.6 Dual Homing to Two BSR Nodes 1193
9.3.20.7 MC Services 1194
9.3.20.8 Routed CO Dual Homing 1196
9.3.20.9 SRRP and Multi-Chassis Synchronization 1202
9.3.20.10 Dual Homing and ANCP 1203
9.3.21 SRRP Enhancement 1203
9.3.21.1 SRRP Fate Sharing 1204
9.3.21.2 Fate Sharing Algorithm 1209
9.3.21.3 SRRP Aware Routing - IPv4/IPv6 Route Advertisement Based on SRRP State 1212
9.3.21.4 SRRP in Conjunction with a PW in ESM Environment – Use Case 1217
9.3.21.5 Group Monitor 1218
9.3.22 Subscriber QoS Overrides 1222
9.3.23 Dual Stack Lite 1225
9.3.23.1 IP-in-IP 1227
9.3.23.2 Configuring Dual Stack Lite 1228
9.3.23.3 L2TP over IPv6 1228
9.3.24 Call Trace 1229
9.4 L2TP Tunnel RADIUS Accounting 1231
9.4.1 Accounting Packets List 1232
9.4.2 RADIUS Attributes Value Considerations 1235
9.4.3 Other Optional RADIUS Attributes 1236

9.4.4 RADIUS VSA to Enable L2TP Tunnel Accounting 1236
9.4.5 MLPPP on the LNS Side 1236
9.5 RADIUS Route Download 1238
9.6 Managed SAP (MSAP) 1240
9.6.1 MSAP QoS Configuration 1246
9.6.2 Sticky MSAP 1247
9.6.3 ESM Identification Process 1248
9.6.3.1 SAP-ID ESM Identifier 1248
9.6.3.2 DSLAM-ID 1248
9.6.4 Default-Subscriber 1249
9.6.5 Multicast Management 1249
9.6.6 Subscriber Mirroring 1249
9.7 Volume and Time Based Accounting 1250
9.7.1 Metering 1250
9.7.1.1 Categories Map and Categories 1250
9.7.1.2 Quota Consumption 1252
9.7.1.3 Minimum Credit Control Quota Values 1252
9.7.1.4 RADIUS VSA Alc-Credit-Control-Quota 1252
9.7.2 Credit Negotiation Mechanisms 1253
9.7.3 Action on Credit Exhaustion 1255
9.7.4 Action on Error-Conditions 1255
9.7.5 Applicability of Volume and Time Based Accounting 1255
9.8 Subscriber Host Idle Timeout 1256
9.9 Web Portal Authentication 1259
9.9.1 HTTP-Redirect (Captive Portal) 1259
9.9.2 One-time HTTP Redirection Overview 1259
9.9.3 Web Authentication Protocol (WPP) 1260
9.9.3.1 WPP Configurations 1261
9.9.3.2 WPP Triggered Host Creation 1262
9.9.3.3 WPP Multi-Chassis Redundancy Support 1263
9.9.3.4 WPP Portal Group 1264
9.10 ESM over MPLS Pseudowires 1265
9.10.1 Encapsulation 1266
9.10.2 ESM and PW Ports 1267
9.10.2.1 ESM on PW-Port Bound to a Physical Port 1268
9.10.2.2 ESM on PXC-Based PW-Ports 1279
9.11 Logical Link Identifier (LLID) 1283
9.12 PADI Authentication Policy for Managed SAP (MSAP) 1284
9.13 Open Authentication Model for DHCP and PPPoE Hosts 1286
9.13.1 Terminology 1286
9.13.2 LUDB and RADIUS Access Models 1286
9.13.3 No Authentication 1287
9.13.4 LUDB Only Access 1287
9.13.5 LUDB Access via DHCPv4 Server 1288
9.13.6 RADIUS Only Access 1288
9.13.7 Consecutive Access to LUDB and RADIUS 1288
9.13.8 RADIUS Fallback 1289
9.14 Flexible Subscriber-Interface Addressing (Unnumbered Subscriber-Interfaces) 1290

9.14.1 Terminology 1290
9.14.2 Flexible Subscriber-Interface Addressing for IPOE/PPPoE v4/v6 Subscribers 1290
9.14.3 Default Gateway in IPv4 Flexible Addressing 1292
9.14.4 IPv4 Subnet Sharing 1293
9.14.5 IPv4 Subnet Mask Auto-Generation 1294
9.14.6 Local-proxy-arp and arp-populate 1295
9.14.7 Gi-address Configuration Consideration 1296
9.14.8 PPPoE Considerations 1296
9.14.9 IPoEv6 Considerations 1297
9.14.10 General Configuration Guidelines for Flexible IP Address Assignment 1297
9.14.11 Caveats 1299
9.15 uRPF for Subscriber Management 1300
9.16 IPoE Sessions 1301
9.16.1 Enabling IPoE Sessions 1302
9.16.2 IPoE Session Authentication 1303
9.16.3 IPoE Session Accounting 1305
9.16.4 IPoE Session Mid-Session Changes 1305
9.16.5 IPoE Session Termination 1306
9.16.6 Limiting the Number of IPoE sessions 1306
9.16.7 SAP Session Index 1306
9.16.8 Resiliency 1307
9.16.9 Notes 1307
9.16.10 Configuration Steps 1308
9.16.11 IPoE Session Migration 1309
9.16.11.1 Additional Notes for IPoE Session Migration of IPv4 Hosts as a Control Channel for Dynamic Data Services 1313
9.17 Data-triggered Subscriber Management 1314
9.17.1 Provisioning Data-triggered ESM 1314
9.17.2 Authentication and Host Creation 1315
9.17.3 DoS Protection 1317
9.17.4 DHCP Promotion 1317
9.17.5 Data-Triggered SLAAC Hosts 1318
9.17.6 Stateful Multi-Chassis Redundancy (MCS) 1319
9.17.7 Stateless Multi-Chassis Redundancy 1319
9.17.7.1 MSAP Support 1321
9.18 RADIUS Subscriber Services 1322
9.18.1 Subscriber Service Building Blocks 1322
9.18.1.1 RADIUS Access-Accept or CoA Message with Subscriber Service Activate or Deactivate VSAs 1323
9.18.1.2 RADIUS Python Interface 1324
9.18.1.3 Python Script 1326
9.18.1.4 Subscriber Service Instance Activation or Deactivation with Optional RADIUS Accounting 1330
9.18.2 Subscriber Services RADIUS VSAs 1332
9.18.3 Subscriber Service RADIUS Accounting 1333
9.18.4 Accounting-Only Subscriber Service 1335
9.18.5 QoS Override-based Subscriber Service 1335

9.18.6 PCC Rule-based Subscriber Services 1338
9.18.6.1 PCC Rule Actions 1339
9.18.6.2 PCC Rule Instantiation 1347
9.18.6.3 PCC Rules in a Subscriber Service 1350
9.18.6.4 Interaction of the PPPoE or IPoE Session QoS Model and PCC Rule-based Subscriber Services 1352
9.18.6.5 PCC Rules on HSMDAv2 1356
9.18.6.6 PCC Rule-based Subscriber Service Activation Failures 1360
9.18.7 Combined Subscriber Services 1361
9.18.8 Subscriber Services Python API 1362
9.18.8.1 Common Subscriber Services Python API 1362
9.18.8.2 Subscriber Service QoS Override Python API 1365
9.18.8.3 Subscriber Service PCC Rules Python API 1366
9.18.9 Operational Commands 1376
9.18.9.1 Show Commands 1376
9.18.9.2 Debug Commands 1379
9.18.9.3 Resource Monitoring 1380
9.19 Residential Gateway Replacement 1381
9.20 ESM Troubleshooting Show Command 1383
9.21 Subscriber Accumulated Statistics 1384
9.22 Hybrid Access 1386
9.22.1 BNG-based HAG 1386
9.22.2 PGW-based HAG 1387
9.23 Connection Bonding 1388
9.23.1 Setup 1388
9.23.2 Downstream Load-Balancing 1389
9.23.3 QoS 1390
9.23.4 Multicast 1391
9.24 Configuring Enhanced Subscriber Management with CLI 1393
9.24.1 Configuring RADIUS Authentication of DHCP Sessions 1393
9.24.2 TCP MSS adjustment for ESM Hosts 1394
9.24.3 Configuring Enhanced Subscriber Management 1394
9.24.3.1 Basic Configurations 1394
9.24.3.2 Subscriber Interface Configuration 1395
9.24.3.3 Configuring Enhanced Subscriber Management Entities 1395
9.24.3.4 Routed CO with Basic Subscriber Management Features 1398
9.24.3.5 Applying the Profiles and Policies 1399
9.24.4 Configuring Dual Homing 1401
9.24.4.1 SHCV Policies 1403
9.24.4.2 SHCV Policy 1403
9.24.4.3 Subscriber Identification Policy 1404
9.24.4.4 Subscriber Profile 1405
9.25 Subscriber Management Command Reference 1407
9.25.1 Command Hierarchies 1407
9.25.1.1 ANCP Commands 1408
9.25.1.2 Authentication Policy Commands 1409
9.25.1.3 BGP Peering Policy Commands 1411
9.25.1.4 Call Trace Commands 1412
9.25.1.5 Category Map and Credit Control Policy Commands 1412

9.25.1.6 Diameter Policy Commands 1416
9.25.1.7 Explicit Subscriber Mapping Commands 1421
9.25.1.8 Filter Commands 1421
9.25.1.9 GSMP Configuration Commands 1422
9.25.1.10 Host Lockout Commands 1422
9.25.1.11 Host Tracking Policy Commands 1424
9.25.1.12 IGMP Policy Commands 1424
9.25.1.13 Multi-Chassis Redundancy Commands 1425
9.25.1.14 PIM Policy Commands 1427
9.25.1.15 RADIUS Accounting Policy Commands 1427
9.25.1.16 RADIUS Route Download Commands 1431
9.25.1.17 Router Advertisement Commands 1431
9.25.1.18 SLA Profile Commands 1434
9.25.1.19 Subscriber Identification Policy Commands 1437
9.25.1.20 Auto-Generated Subscriber Identification Key Commands 1438
9.25.1.21 Auto-Generated Subscriber Identification Key Service Commands 1438
9.25.1.22 Subscriber MCAC Policy Commands 1439
9.25.1.23 Subscriber Profile Commands 1439
9.25.1.24 IPoE Session Policy Commands 1441
9.25.1.25 SHCV Policy Commands 1442
9.25.1.26 Subscriber Management Service Commands 1443
9.25.1.27 Show Commands 1472
9.25.1.28 Monitor Commands 1475
9.25.1.29 Clear Commands 1476
9.25.1.30 Debug Commands 1477
9.25.1.31 Tools Commands 1479
9.25.2 Triple Play Subscriber Management Configuration Commands 1482
9.25.2.1 Triple Play Subscriber Management Configuration Commands 1482
9.25.2.2 Show Commands 1921
9.25.2.3 Clear Commands 2036
9.25.2.4 Tools Commands 2046
9.25.2.5 Debug Commands 2070
9.25.2.6 Monitor Commands 2089

10 Oversubscribed Multi-Chassis Redundancy (OMCR) in ESM 2095
10.1 Overview 2095
10.1.1 Terminology and Abbreviations 2095
10.1.2 Restrictions 2095
10.2 Deploying Oversubscribed Multi-Chassis Redundancy 2097
10.2.1 Resource Exhaustion Notification and Simultaneous Failures 2100
10.2.2 Resource Monitoring 2100
10.2.3 Warm-Standby Mode Of Operation 2102
10.2.4 IPoE vs PPPoE 2104
10.2.5 Persistency 2104
10.2.6 Routing and Redundant Interface in OMCR 2104
10.2.7 Revertive Behavior 2107
10.2.8 Service Restoration Times 2107
10.2.9 Processing of the SRRP Flaps 2108

10.2.10 Accounting 2108
10.2.11 Configuration Guidelines 2109
10.2.12 Troubleshooting Commands 2110
10.3 OMCR Command Reference 2117
10.3.1 Command Hierarchies 2117
10.3.2 OMCR Configuration Commands 2117
10.3.2.1 OMCR Configuration Commands 2117

11 WIFI Aggregation and Offload 2119
11.1 WIFI Aggregation and Offload Overview 2119
11.2 Layer 2 over Soft-GRE Tunnels 2120
11.2.1 Encapsulation 2121
11.2.2 Data Path 2125
11.2.3 MDA-based Redundancy 2126
11.3 Tunnel Level Egress QoS 2127
11.3.1 QoS Overrides 2130
11.3.2 Operational Commands 2131
11.4 Authentication 2136
11.4.1 EAP-Based Authentication 2137
11.4.1.1 RADIUS Proxy 2138
11.4.2 Portal Authentication 2142
11.4.3 AA-based Portal Redirection 2144
11.5 Address Assignment 2145
11.6 WIFI Mobility Anchor 2147
11.7 Wholesale 2148
11.8 CGN on WLAN-GW 2149
11.9 Lawful Intercept on WLAN-GW 2149
11.10 WLAN Location Enhancements 2150
11.10.1 Triggered Interim Accounting-Updates 2150
11.10.2 Mobility Triggered Interim Updates with Counters 2152
11.10.3 Operational Support 2153
11.11 3G/4G Interworking 2154
11.11.1 Signaling Call Flow 2154
11.11.1.1 GTP Setup with EAP Authentication 2154
11.11.2 Location Notification in S2a 2155
11.11.2.1 WLAN Location over S2a 2156
11.11.2.2 Cellular Location over S2a 2156
11.11.2.3 Cellular Location over Gn Interface 2157
11.11.2.4 Operational Support 2157
11.12 Migrant User Support 2159
11.12.1 Migrant User Support with Portal-Authentication 2159
11.12.1.1 DHCP 2159
11.12.1.2 Authentication and Forwarding 2160
11.12.2 Migrant User Support with EAP Authentication 2161
11.12.3 Data Triggered Subscriber Creation 2161
11.13 Distributed Subscriber Management (DSM) 2166
11.13.1 DHCP 2167
11.13.2 Authentication and Accounting 2167
11.13.2.1 DSM Data-Plane 2170

11.13.3 IP Filtering 2170
11.13.3.1 Policing 2171
11.13.3.2 Lawful Intercept (LI) 2172
11.13.3.3 Data-Triggered UE Creation 2173
11.13.4 Idle-Timeout and Session-Timeout Management 2173
11.13.5 Operational Commands 2174
11.13.6 Pool Manager 2175
11.13.7 DHCPv6 and SLAAC 2176
11.13.8 Application Assurance Support 2178
11.14 Call Trace 2178
11.15 Distributed RADIUS Proxy 2179
11.15.1 Enhanced Subscriber Management 2181
11.15.2 Distributed Subscriber Management 2181
11.15.3 VLAN Awareness 2181
11.15.4 Operational Commands 2182
11.16 WLAN-GW 1:1 Active-Backup Redundancy 2184
11.16.1 DHCP Server Redundancy 2185
11.16.2 Subscriber Creation after Switchover 2186
11.17 WLAN-GW Triggered Stateless Redundancy (N:1) 2187
11.18 AP Triggered Stateless WLAN-GW Redundancy (N:1) 2187
11.19 IPv6-only Access 2188
11.19.1 IPv6 GRE Tunnels 2188
11.19.2 IPv6 Client-Side RADIUS Proxy 2189
11.19.3 Dual-Stack UEs over WLAN-GW 2190
11.19.3.1 SLAAC Prefix Assignment 2190
11.19.3.2 DHCPv6 IA_NA Assignment 2191
11.19.3.3 Migrant User Support 2191
11.19.3.4 Accounting 2192
11.20 Layer 2 Wholesale 2194
11.21 VLAN to WLAN-GW IOM/IMM Steering via Internal Epipe 2195
11.22 Soft-L2TPv3 Tunnels 2197
11.23 WLAN-GW — Dynamic Tunnel X-Connect for Seamless Inter-WLAN-GW Mobility 2199
11.23.1 Processing on the V-GW 2200
11.23.2 Processing on H-GW 2200
11.23.3 Idle Timeout Handling 2201
11.23.4 Distributed RADIUS Proxy for Closed SSID 2201
11.23.5 H-GW Redundancy 2201
11.24 ISA Operational Commands and Key Performance Indicators 2202
11.24.1 ISA Resources 2202
11.24.2 ISA Load 2202
11.24.3 Packet Statistics 2203
11.25 WiFi Command Reference 2205
11.25.1 Command Hierarchies 2205
11.25.1.1 WLAN-GW Commands 2205
11.25.1.2 ISA Commands 2206
11.25.1.3 WLAN-GW Service Commands 2208
11.25.1.4 Data Plane Related Commands 2212
11.25.1.5 RADIUS Server and Proxy Commands 2213

11.25.1.6 LUDB Matching for RADIUS Proxy Cache 2216
11.25.1.7 Port Policy Commands 2216
11.25.1.8 WIFI Aggregation and Offload – Migrant User Support Commands 2217
11.25.1.9 Show Commands 2217
11.25.1.10 Debug Commands 2218
11.25.1.11 Tools Commands 2219
11.25.1.12 Clear Commands 2219
11.25.2 WIFI Aggregation and Offload Commands 2220
11.25.2.1 WIFI Aggregation and Offload Commands 2220
11.25.2.2 CLI Command Description for RADIUS Server 2240
11.25.2.3 CLI Command Description for RADIUS Proxy Server 2243
11.25.2.4 LUDB Matching of RADIUS Proxy Cache Commands 2251
11.25.2.5 WLAN-GW-Group Commands 2253
11.25.2.6 Port Policy Commands 2260
11.25.2.7 WLAN-GW Group Interface Commands 2261
11.25.2.8 Migrant User Support Commands 2288

12 GTP 2353
12.1 GTP Uplink 2353
12.1.1 Identification Attributes 2353
12.1.2 P-GW/GGSN Selection 2353
12.1.3 Configuration 2354
12.1.4 QoS Support 2355
12.1.5 GTP Session Hold 2356
12.1.6 Selective Breakout 2357
12.1.7 IPoE Support 2358
12.1.8 PPPoE Support 2359
12.2 GTP Access 2360
12.2.1 GTP Termination 2360
12.2.1.1 Multiple APNs 2362
12.2.2 GTP Session Setup 2362
12.2.2.1 Supported IP Stacks 2363
12.2.3 QoS 2363
12.2.4 Multicast 2364
12.3 DHCP over GTP-u 2365
12.3.1 Address Management Related PCOs 2365
12.3.2 Address Allocation Modes 2366
12.4 GTP Peering 2368
12.5 GTP Command Reference 2369
12.5.1 Command Hierarchies 2369
12.5.1.1 GTP Commands 2369
12.5.1.2 GTP Service Commands 2371
12.5.1.3 GTP Show Commands 2372
12.5.1.4 GTP Clear Commands 2372
12.5.1.5 GTP Debug Commands 2373
12.5.2 GTP Command Descriptions 2373
12.5.2.1 GTP Commands 2373

13 Virtual Residential Gateway 2405
13.1 Overview 2405
13.1.1 Access Modes 2406
13.1.2 Home Context on the vRGW 2407
13.1.2.1 Implicit Home Authentication 2407
13.1.2.2 Explicit Home Authentication 2409
13.1.2.3 Change of Configuration 2410
13.1.2.4 Home Lifetime 2410
13.1.3 Device Context on the vRGW 2411
13.1.4 Dynamic Configuration Changes 2411
13.1.5 Per-Home Pool Management and L2-Aware NAT 2412
13.1.5.1 Sticky IP Addresses 2412
13.1.5.2 Managed Static IPv4 Addresses 2413
13.1.5.3 DMZ 2413
13.1.6 IPv6 2413
13.1.7 QoS and Filter Support 2414
13.1.8 Data-Triggered Authentication 2415
13.1.9 Per-Host NAT Port Ranges 2416
13.1.10 Inter-Chassis Redundancy 2417
13.1.10.1 Pool State Synchronization 2420
13.1.10.2 Regular Group Interfaces 2421
13.1.10.3 WLAN-GW Group Interfaces 2421
13.1.11 BRG and vRG Caveats 2422
13.1.12 External Allocation of L2-Aware NAT Outside IP Addresses 2423
13.1.13 PPPoE Client 2426
13.1.13.1 PPPoE Client Setup 2426
13.1.13.2 PPPoE Client Failure 2427
13.1.13.3 LCP Keepalive 2427
13.1.13.4 MRU/MTU 2428
13.1.14 SLAAC Prefix Replacement 2428
13.2 Home LAN Extension 2429
13.2.1 Overview 2429
13.2.2 Authentication and Authorization 2430
13.2.3 Data Plane Tables 2430
13.2.4 BGP EVPN VPLS 2431
13.2.5 Assistive Address Resolution 2432
13.2.6 MAC Address Translation 2433
13.2.7 Configuring HLE 2434
13.3 AP Agnostic Access for Multiple Dwelling Units 2436
13.3.1 Overview 2436
13.3.2 Bridge Domain and BRG Identification 2437
13.3.3 ARP Handling 2438
13.3.4 Mobility 2438
13.4 vRGW and HLE Command Reference 2439
13.4.1 Command Hierarchies 2439
13.4.1.1 Subscriber Management BRG Profile Commands 2439
13.4.1.2 Subscriber Management Commands 2440
13.4.1.3 Service Commands 2441
13.4.1.4 Router HLE Commands 2441

13.4.1.5 Subscriber Management HLE Commands 2442
13.4.1.6 Service Commands 2442
13.4.1.7 Show Commands 2443
13.4.1.8 Debug Commands 2443
13.4.1.9 Clear Commands 2444
13.4.1.10 Tools Commands 2444
13.4.2 Command Descriptions 2445
13.4.2.1 Generic Commands 2445
13.4.2.2 vRGW Commands 2446
13.4.2.3 Subscriber Management Commands 2454
13.4.2.4 Service Commands 2459
13.4.2.5 Router Home LAN Extension Commands 2464
13.4.2.6 Subscriber Management Commands 2466
13.4.2.7 Service Commands 2469
13.4.2.8 Show Commands 2473
13.4.2.9 Debug Commands 2482
13.4.2.10 Clear Commands 2484
13.4.2.11 Tools Commands 2485

14 Service Chaining for ESM Hosts with L2-Aware NAT 2491
14.1 Steering to Service Chains for ESM hosts with L2-Aware NAT 2491
14.1.1 Terminology 2492
14.2 VAS Filters on the ISA 2493
14.2.1 Matching 2493
14.2.2 Forwarding 2493
14.2.3 NSH Insertion 2494
14.2.4 Configuration 2494
14.3 EVPN Route Updates and Tracking 2496
14.3.1 NVE Bridging to SF 2496
14.3.2 NVE Routing to SF 2497
14.4 Data Path on the Subscriber Edge 2500
14.4.1 Upstream Traffic (Access to Network) 2500
14.4.2 Downstream Traffic — From Network 2501
14.5 Data Path on NVE 2501
14.6 Service Chaining for ESM Hosts with L2-Aware NAT 2503
14.6.1 Command Hierarchies 2503
14.6.1.1 ISA Service Chaining Commands 2503
14.6.1.2 Show Commands 2504
14.6.2 Command Descriptions 2505
14.6.2.1 Generic Commands 2505
14.6.2.2 Service Chaining Commands 2506
14.6.2.3 Show Commands 2516

15 Dynamic Data Services 2521
15.1 Introduction to Dynamic Data Services 2521
15.2 RADIUS-Triggered Dynamic Data Services Associated With a PPPoE or IPoE Session as Control Channel 2522
15.3 Data-Triggered Dynamic Data Services 2522
15.3.1 Data Trigger 2523

15.3.2 Dynamic Services Data Trigger Capture SAP 2524
15.3.3 RADIUS Authentication 2527
15.3.4 Local Authentication 2529
15.3.5 Data-Triggered Dynamic Service Provisioning 2531
15.3.6 Control Plane Protection 2535
15.3.7 Debugging 2539
15.4 Dynamic Data Services Python API 2540
15.5 RADIUS Triggered Dynamic Data Services Command Reference 2543
15.5.1 Command Hierarchies 2543
15.5.1.1 RADIUS Triggered Dynamic Data Services Commands 2543
15.5.1.2 Basic System Command 2545
15.5.1.3 Show Commands 2545
15.5.1.4 Clear Commands 2545
15.5.1.5 Debug Commands 2546
15.5.1.6 Tools Commands 2546
15.5.2 RADIUS Triggered Dynamic Data Services Commands 2547
15.5.2.1 Configuration Commands 2547
15.5.2.2 Basic System Command 2561
15.5.2.3 Show Commands 2562
15.5.2.4 Clear Commands 2575
15.5.2.5 Debug Commands 2575
15.5.2.6 Tools Commands 2578

16 Diameter and Diameter Applications 2581
16.1 Restrictions 2581
16.2 Terminology 2583
16.3 Diameter Base 2585
16.3.1 Diameter Base Protocol 2585
16.3.2 Diameter Peers and the Role of a Diameter Node in SR OS 2586
16.3.3 Capability Exchange 2586
16.3.4 Connection Termination 2587
16.3.5 Diameter Hosts and Realms 2588
16.3.5.1 Forwarding/Routing of Diameter Messages 2588
16.3.5.2 Configuration of Host Names and Realms in SR OS 2589
16.3.6 Dynamically Learned Parameters 2591
16.3.7 Retransmissions and Message Timers 2591
16.3.8 Python Support 2593
16.4 3GPP-Based Diameter Credit Control Application (DCCA) - Online Charging 2594
16.4.1 Diameter Gy Out Of Credit Actions 2600
16.4.1.1 Graceful Service Termination 2600
16.4.2 Extended Failure Handling (EFH) 2602
16.4.2.1 Extended Failure Handling Example Call Flow 2604
16.4.2.2 Extended Failure Handling Triggers 2606
16.4.2.3 Assigning Interim Credit 2608
16.4.2.4 Enabling Extended Failure Handling 2610
16.4.2.5 Configuration Example 1 - Single Volume Interim Credit Value 2611
16.4.2.6 Configuration Example 2 - Interim Credit Values Per Rating Group 2611
16.4.2.7 Monitoring the Extended Failure Handling State 2613

16.4.2.8 Additional Call Flow Examples 2616
16.4.3 Gy CCR-t Replay 2621
16.5 Policy Management via Gx Interface 2624
16.5.1 Gx Protocol 2625
16.5.2 Policy Assignment Models 2626
16.5.3 IP-CAN Session – Gx Session Identification 2628
16.5.3.1 User Identification in PCRF 2629
16.5.3.2 NAS-Port-Id as Subscription-Id 2631
16.5.4 Gx Interface and ESM Subscriber Instantiation 2631
16.5.4.1 Gx and Dual-Stack Hosts 2635
16.5.4.2 Gx and PPPoEv6-DHCP 2637
16.5.5 Gx Fallback Function 2638
16.5.6 Gx CCR-I Replays 2640
16.5.7 Gx CCR-t Replays 2640
16.5.7.1 RAR and CCR-t Replay 2641
16.5.7.2 CCR-t Replay And Multi-Chassis Redundancy 2641
16.5.7.3 CCR-t Replay And High Availability 2642
16.5.8 Automatic Updates for IP Address Allocation/De-allocation 2642
16.5.9 DHCPv4/v6 Re-Authentication and RADIUS CoA Interactions With Gx 2643
16.5.10 Gx, ESM and AA 2644
16.5.10.1 ESM Subscriber-Host vs AA Subscriber 2645
16.5.10.2 AA Subscriber State 2645
16.5.11 Policy Management via Gx 2646
16.5.12 Gx-Based Overrides 2646
16.5.12.1 Instantiation of Gx Overrides 2648
16.5.12.2 HTTP Redirect Override 2651
16.5.12.3 Removal of Overrides 2653
16.5.12.4 Examples of Gx Overrides 2653
16.5.13 PCC Rules 2658
16.5.13.1 PCC Rule Concept 2659
16.5.13.2 PCC Rule Instantiation 2660
16.5.13.3 Base QoS-Policy and Base Filter 2663
16.5.13.4 Generic Policy Sharing and Rule Sharing 2664
16.5.13.5 PCC Rule Name and PCC Rule Removal 2665
16.5.13.6 Gx Rule Ordering 2666
16.5.13.7 PCC Rule Override 2668
16.5.13.8 Aggregation of IP-Criterion 2669
16.5.13.9 Combining IPv4 and IPv6 Entries within the Rule 2669
16.5.13.10 Gx Rules with Multiple Actions and Action Sharing 2670
16.5.13.11 Alc-NAS-Filter-Rule-Shared AVP vs Flow-Information AVP 2670
16.5.13.12 RADIUS and Gx Interaction 2671
16.5.13.13 Bulk Changes via CLI while Gx Rules are Active 2672
16.5.13.14 PCC Rule Direction 2674
16.5.13.15 Action 2674
16.5.13.16 Rate-Limiting Action (Ingress, Egress) 2675
16.5.13.17 Forwarding-Class Change (Ingress, Egress) 2679
16.5.13.18 QoS Forward (Ingress and Egress) 2680
16.5.13.19 Next Hop Redirect (Ingress) 2681

16.5.13.20 HTTP Redirect (Ingress) 2681
16.5.13.21 Filter Forward/Drop (Ingress and Egress) 2682
16.5.13.22 Service Gating Function 2682
16.5.13.23 PCC Rule Provisioning Example 2683
16.5.13.24 Operational Aspects 2684
16.5.13.25 PCC Rules and Capacity Planning 2685
16.5.13.26 PCC Rule Scaling Example 2686
16.5.14 NAS Filter Inserts 2687
16.5.14.1 Examples of NAS Entry Inserts 2688
16.5.15 Error Handing and Rule Failure Reporting in ESM 2688
16.5.15.1 AVP Decoding Failure in Gx 2689
16.5.15.2 ESM Rule-Installation Failure 2690
16.5.15.3 Failure Reporting in AA 2692
16.5.15.4 Summary of Failure Reporting 2692
16.5.16 Usage-Monitoring and Reporting 2694
16.5.16.1 ESM Usage-Monitoring - What is Being Monitored 2694
16.5.16.2 AA Usage-Monitoring – What is Being Monitored 2696
16.5.16.3 Requesting Usage-Monitoring in ESM 2696
16.5.16.4 Reporting Accumulated Usage 2697
16.5.16.5 Disabling Usage-Monitoring 2698
16.5.16.6 Usage-Monitoring for PCC Rules 2698
16.5.16.7 Session Termination 2699
16.5.16.8 Usage Monitoring When Multiple Subscriber Hosts or Sessions Share an SLA Profile Instance 2699
16.5.16.9 Usage-Monitoring Examples 2700
16.5.17 Event Triggers 2702
16.5.18 Subscriber Verification 2703
16.5.19 Subscriber Termination 2703
16.5.20 Mobility Support in WiFi 2704
16.5.20.1 Redundancy 2704
16.5.21 Persistency and Origin-State-ID AVP 2704
16.5.22 Overload Protection 2705
16.6 Diameter NASREQ Application 2706
16.6.1 Sample Configuration Steps 2709
16.7 Diameter Redundancy 2712
16.7.1 Diameter Peer Level Redundancy 2712
16.7.2 Diameter Multi-Chassis Redundancy 2713
16.7.2.1 Diameter Proxy Model General Operational Principles 2714
16.7.2.2 Diameter Proxy Activity Selection 2717
16.7.2.3 Synchronization and MCS 2717
16.7.2.4 Retransmissions 2718
16.7.2.5 Retransmissions and the T-bit 2722
16.7.2.6 Diameter Proxy Role 2722
16.7.2.7 Diameter Proxy and CC-Request-Number AVP 2724
16.7.2.8 Stateless Diameter Proxy 2724
16.7.2.9 Switchover Scenarios 2724
16.7.2.10 Log/Trap Generation Caused by Diameter Proxy State Change 2728
16.7.2.11 Switchover Update Event (CCR-u) 2728
16.7.2.12 Isolated Chassis 2728

16.7.2.13 Diameter Identities 2729
16.7.2.14 High Availability 2729
16.7.3 Gx Specific Behavior 2729
16.8 Diameter Debugging....2731

17 Python Script Support for ESM 2737
17.1 Python Script Support for ESM....2737
17.2 Python in SR OS Overview 2737
17.2.1 Python Policy – GTPv1-C API 2737
17.2.2 Python Policy – GTPv2-C API 2739
17.2.3 Python Changes 2741
17.3 Python Support in sub-ident-policy....2742
17.3.1 Configuration 2743
17.3.2 Operator Debugging 2744
17.3.3 Python Scripts 2745
17.3.4 Sample Python Scripts 2746
17.3.4.1 Example....2746
17.3.4.2 Example....2747
17.3.4.3 Example....2748
17.3.5 Limitations 2750
17.4 RADIUS Script Policy Overview 2751
17.4.1 Python RADIUS API 2752
17.4.2 Sample Script 2754
17.5 Python Policy Overview....2755
17.5.1 Python Policy – RADIUS API 2756
17.5.2 Python Policy – DHCPv4 API 2756
17.5.3 Python Policy – DHCPv6 API 2759
17.5.4 Python Policy – Diameter API 2766
17.5.5 Python Policy – DHCP Transaction Cache API....2772
17.5.6 Python for PPPoE API....2774
17.5.7 Python API for PPP Packet 2776
17.5.8 Python API for PPP PAP 2777
17.5.9 Python API for PPP CHAP 2778
17.5.10 Python ESM API....2779
17.5.11 Python Cache Support 2781
17.5.12 Applying a Python Policy 2782
17.5.13 Python Script Protection 2783
17.6 Tips and Tricks 2784
17.7 Python Command Reference 2785
17.7.1 Command Hierarchies....2785
17.7.1.1 Python Policy Commands 2785
17.7.1.2 Python Script Commands....2786
17.7.1.3 Services Commands 2787
17.7.1.4 Show Commands 2788
17.7.1.5 Tools Commands 2788
17.7.1.6 Debug Commands....2789
17.7.1.7 Clear Commands....2789
17.7.2 Command Descriptions 2789
17.7.2.1 Python Configuration Commands....2789

17.7.2.2 Show Commands 2808
17.7.2.3 Debug Commands 2810
17.7.2.4 Clear Commands 2812
17.8 Python RADIUS Command Reference 2813
17.8.1 Command Hierarchies 2813
17.8.2 Command Descriptions 2814
17.8.2.1 Generic Commands 2814
17.8.2.2 Script Commands 2815

18 Standards and Protocol Support 2821

1 Getting Started

1.1 About This Guide

This guide describes details pertaining to Triple Play Services Delivery Architecture (TPSDA) support provided by the operating system and presents examples to configure and implement various protocols and services.

This guide is organized into functional chapters and provides concepts and descriptions of the implementation flow, as well as Command Line Interface (CLI) syntax and command usage.

The topics and commands described in this document apply to the:

- 7450 ESS
- 7750 SR
- VSR

Table 1 lists the available chassis types for each SR OS router.

Table 1 **Supported SR OS Router Chassis Types**

7450 ESS	7750 SR
• 7450 ESS-7/12 running in standard mode (not mixed-mode)	• 7450 ESS-7/12 running in mixed-mode (not standard mode) • 7750 SR-a4/a8 • 7750 SR-c4/c12 • 7750 SR-1e/2e/3e • 7750 SR-7/12 • 7750 SR-12e • 7750 SR-7s/12s • 7750 SR-1

For a list of unsupported features by platform and chassis, refer to the SR OS 16.0.R*x* Software Release Notes, part number 3HE 14220 000*x* TQZZA or the *VSR Release Notes*, part number 3HE 14204 000*x* TQZZA.

Command outputs shown in this guide are examples only; actual displays may differ depending on supported functionality and user configuration.

Note: This guide generically covers Release 16.0.R*x* content and may contain some content that is released in later maintenance loads. Refer to the SR OS 16.0.R*x* Software Release Notes, part number 3HE 14220 000*x* TQZZA or the *VSR Release Notes*, part number 3HE 14204 000*x* TQZZA, for information about features supported in each load of the Release 16.0.R*x* software.

1.2 Nokia SR OS Services Configuration Process

Table 2 lists the tasks necessary to configure TPSDA entities. Each section describes a software area and provides CLI syntax and command usage to configure parameters for a functional area.

Table 2 **Configuration Process**

Area	Task	Section
Triple play services	Provision and deploy services	Deploying and Provisioning Services
	Configure triple play services	Configuring Triple Play Services with CLI
Protocol configuration	Configure DHCP	Configuring DHCP with CLI
	Configure traffic steering on L2TP LAC	Configuring Traffic Steering on L2TP LAC
PPPoE	Configure point-to-point over Ethernet	Point-to-Point Protocol over Ethernet Management
Security	Configure triple play security	Configuring Triple Play Security with CLI
Multicast	Configure triple play multicast services	Configuring Triple Play Multicast Services with CLI
Enhanced Subscriber management	Configure RADIUS server	Uniform RADIUS Server Configuration
	Use scripts dynamic recognition of subscribers	Using Scripts for Dynamic Recognition of Subscribers
	Configure IP and IPv6 filter policies for subscriber hosts	Configuring IP and IPv6 Filter Policies for Subscriber Hosts
	Configure dual stack lite	Dual Stack Lite
	Configure MSAP	MSAP QoS Configuration
	Configure web authentication protocol (WPP)	WPP Configurations
	Configure IPoE session	Configuration Steps
	Troubleshoot ESM	ESM Troubleshooting Show Command
	Configure enhanced subscriber management	Configuring Enhanced Subscriber Management with CLI
	Deploy oversubscribed multi-chassis redundancy	Deploying Oversubscribed Multi-Chassis Redundancy
	Python Scripting	Python Script Support for ESM

2 Introduction to Triple Play

2.1 Nokia's Triple Play Service Delivery Architecture

This section discusses Nokia's Triple Play Service Deliver Architecture (TPSDA) implementation.

2.1.1 Introduction to Triple Play

For more than a decade, telephony service providers have considered offering video services to residential customers. However, in the past it was not economically nor technically feasible to launch the implementation on a large scale.

Several technical trends and evolutions have propelled video delivery to the foreground, including:

- Technical improvements in areas such as real-time MPEG encoding and compression.
- Widespread deployment of High Speed Internet (HSI) over broadband access (ADSL and cable modems).
- Decreased cost of high-bandwidth infrastructure (typically Ethernet-based) as well as storing, converting, and delivering video content.
- Increased competition between telephony and cable operators. This is partly due to changes in regulations.

Traditional cable operators began offering television services and later added Internet access and telephony to their offerings. Conversely, traditional telephony operators such as RBOCs, PTTs, have also added Internet access, and many are now in the process of also adding video delivery.

This bundling of video, voice, and data services to residential subscribers is now commonly known as Triple Play services. The video component always includes linear programming (broadcast television), but often also has a non-linear Video on Demand (VoD) component.

2.1.2 Blueprint for Optimizing Triple Play Service Infrastructures

Nokia's TPSDA allows network operators to progressively integrate their HSI, voice, and video services within a unified and homogeneous Ethernet-based aggregation network environment. The key benefits of the proposed service infrastructure include cost optimization, reduced risk, and accelerated time to market for new services.

At a high level, TPSDA implements:

- Ethernet-based service architecture — Solves bandwidth bottlenecks and exponential capital expenditure and operating expenses issues in the second mile by leveraging the efficiency of this technology.
- Multiple distributed service edges — Allows service providers to achieve faster times to market for new services while retaining the existing Broadband Remote Access Server (BRAS) / and Point-to-Point Protocol over Ethernet (PPPoE) mode of operation for wholesale and retail HSI.
- Distributed multicasting functions in access and aggregation networks — Enables service providers to optimize bandwidth and content delivery mechanisms, based on densities and penetration rates. It is also essential to subscriber and service scaling, and optimizes the bandwidth required in the aggregation network.
- Carrier video and Voice over Internet Protocol (VoIP) services using Dynamic Host Configuration Protocol (DHCP) — Enables service providers to introduce plug-and-play services delivered through set-top boxes and VoIP devices, which are designed for use with the DHCP.
- Flexible deployment models — The architecture allows data, video, and VoIP services to be rapidly rolled out without any lock-in to specific operational models. It allows service providers to maximize flexibility and minimize financial and technological risks by allowing all modes of operation, including:
  - Copper (DSL/DSLAM) and fiber-based (FTTx) deployments in the first mile.
  - Single or multiple last mile circuits.
  - Bridged or routed home gateways.
  - Single or multiple IP address deployment models.

2.1.3 Architectural Foundations

With the SR OS, the architectural foundations of Nokia's TPSDA is reinforced while its applicability is expanded to encompass many new deployment models and support Any Mode of Operation (AMO). Through these enhancements, TPSDA becomes more universally deployable and flexible in addressing the specifics of any provider's network rollout.

Nokia has defined new terminologies that have been adopted industry-wide, including:

- Broadband Service Access Node (BSAN)
- Broadband Service Aggregator (BSA)
- Broadband Service Router (BSR)

Figure 1 depicts TPSDA's centralized, integrated element, service and subscriber management architecture.

Figure 1 **Triple Play Service Delivery Architecture**

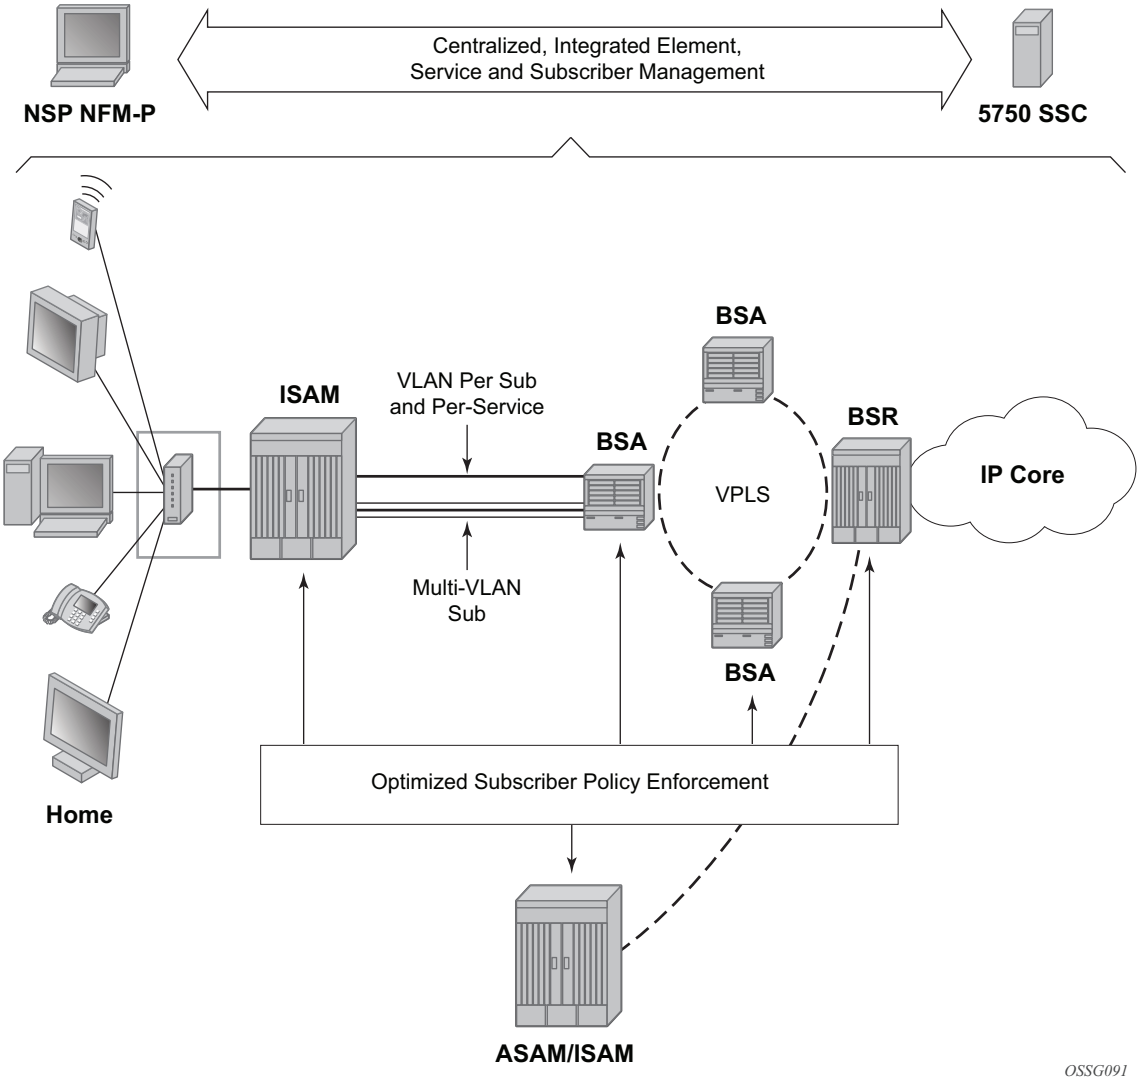

2.1.4 Optimizing Triple Play Service Infrastructures

More than a branding exercise, new terminologies signal a significant shift from “best effort” and traditional DSLAMs, Ethernet switches and BRASs, in the sense that they capture a shift in required characteristics and capabilities for a new generation of service rollouts, including:

- High-availability for non-stop service delivery (non-stop unicast and multicast routing, non-stop services, and so on).

- Multi-dimensional scale (such as the ability to scale performance, bandwidth, services, and subscribers concurrently).
- Ethernet — Optimization (leading density, capacity, scaling, performance).
- Optimal system characteristics (optimal delay, jitter, and loss characteristics, and so on).
- Rich service capabilities with uncompromised performance.

Nokia's Triple Play Service Delivery Architecture (TPSDA) advocates the optimal distribution of service intelligence over the BSAN, BSA and BSR, rather than concentrating on fully centralized or decentralized BRAS models which artificially define arbitrary policy enforcement points in the network. With the SR OS, the optimized enforcement of subscriber policies across nodes or over a single node (as dictated by evolving traffic patterns), allows a more flexible, optimized, and cost-effective deployment of services in a network, guaranteeing high quality and reliable delivery of all services to the user. Nokia's TPSDA entities are described in Table 3.

Table 3 **Nokia's TPSDA**

Entity	Description
Subscriber Management	Centralized and fully integrated with element and services management across the infrastructure end-to-end solution (through the Nokia 5750 SSC).
Policy Enforcement	Optimally distributed, based on actual traffic patterns. Maximized flexibility, minimized risk of architectural lock-in. Optimized cost structure.
Support for "Any Mode of Operation"	With TPSDA, network economics, subscriber density, network topologies and subscriber viewership patterns define the optimal policy enforcement point for each policy type (security, QoS, multicasting, anti-spoofing, filtering, and so on). The SR OS capabilities allow service providers to support any mode of operation, including any combination of access methods, home gateway type, and policy enforcement point (BSAN, BSA or BSR or a combination of the three).

All of the SR OS and Nokia's 5750 SSC's subscriber policy enforcement and management capabilities described in this section build upon Nokia's TPSDA extensive capabilities and provide key capabilities in the following areas:

- Operationalization of Triple Play Services (AAA, subscriber policy enforcement, and so on)
- Service Assurance and Control
- Non-stop Video Service Delivery

2.1.4.1 Distributed Service Edges

The TPSDA architecture (Figure 2), is based on two major network elements optimized for their respective roles, the Broadband Service Aggregator (BSA) and the Broadband Service Router (BSR). An important characteristic of BSAs and BSRs is that they effectively form a distributed virtual node with the BSAs performing subscriber-specific functions where the various functions scale, and the BSRs providing the routing intelligence where it is most cost-effective.

The Nokia 7450 ESS and 7750 SR OS, respectively, provide the BSA and BSR functionalities in TPSDA. Both are managed as a single virtual node using Nokia's NSP NFM-P, which provides a unified interface for streamlined service and policy activation across the distributed elements of the TPSDA architecture, including VPLS, QoS, multicasting, security, filtering, and accounting.

Figure 2 **Nokia's Triple Play Service Delivery Architecture**

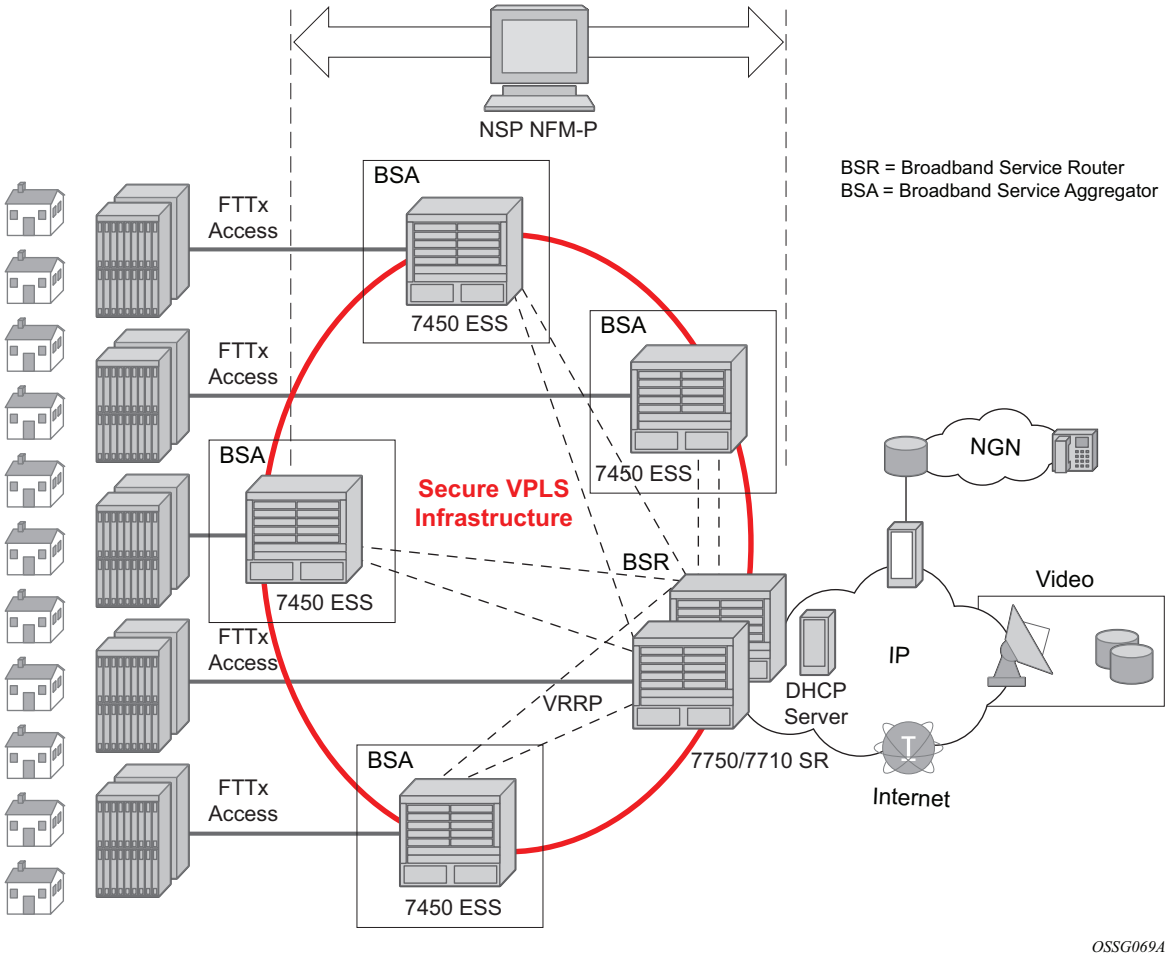

Digital subscriber line access multiplexers (DSLAMs) or other access nodes are connected to Ethernet access ports on the BSA. Typically a single VLAN per subscriber is configured between the access node and the BSA. A VLAN per subscriber provides a persistent context against which per-subscriber policies (QoS, filtering, accounting) can be applied in the BSA.

Scaling of traffic and services is achieved by dividing the Layer 2 and Layer 3 functions between the BSA and BSR and by distributing key service delivery functions. BSAs are more distributed than BSRs, cost-effectively scaling per-subscriber policy enforcement.

The BSA is a high-capacity Ethernet-centric aggregation device that supports hundreds of gigabit Ethernet ports, tens of thousands of filter policies, and tens of thousands of queues. The BSA incorporates wire speed security, per-subscriber service queuing, scheduling, accounting, and filtering.

BSAs aggregate traffic for all services towards the BSR. The BSR terminates the Layer 2 access and routes over IP/MPLS (Multi Protocol Label Switching) with support for a full set of MPLS and IP routing protocols, including multicast routing. The BSR supports hundreds of ports and sophisticated QoS for per-service and per-content/source differentiation.

The connectivity between BSAs and BSRs is a Layer 2 forwarding model shown in Figure 2 above as a secure VPLS infrastructure. This refers to the fact that the BSA-BSR interconnections form a multipoint Ethernet network with security extensions to prevent unauthorized communication, denial of service, and theft of service. One of the advantages of using VPLS for this application is that VPLS instances can be automatically established over both ‘hub and spoke’ and ring topologies providing sub-50 ms resilience. Regardless of the fiber plant layout, VPLS enables a full mesh to be created between BSA and BSR nodes, ensuring efficient traffic distribution and resilience to node or fiber failure.

Other unique features of the BSA and BSR that contribute to this secure VPLS infrastructure are:

1. Using Residential Split Horizon Groups (RSHG), direct user-user bridging is automatically prohibited, without the need for address-specific ACLs.
2. RSHG combined with the ARP reply agent perform ARP and broadcast suppression to ensure that addressing information is restricted.
3. Protection against theft of service and denial of service is provided by MAC or IP filters automatically populated using DHCP snooping, and by MAC pinning.
4. Using the RADIUS interface, is possible to perform RADIUS authentication of users before allowing a DHCP discover to progress into the network.

2.1.4.2 Service Differentiation, QoS Enablement

Nokia's TPSDA approach provides a model based on call admission for video and VoIP, with the need to guarantee delay/jitter/loss characteristics once the service connection is accepted. The architecture also meets the different QoS needs of HSI, namely per-subscriber bandwidth controls, including shaping and policing functions that have little or no value for video and VoIP services. In conjunction with the architecture's support for content differentiation, this enables differentiated service pricing within HSI.

The distribution of QoS policy and enforcement across BSA and BSR allows the service provider to implement meaningful per-subscriber service level controls. Sophisticated and granular QoS in the BSA allows the service provider to deliver truly differentiated IP services based on the subscriber as well as on the content.

In the BSR to BSA downstream direction (Figure 3), IP services rely on IP layer classification of traffic from the network to queue traffic appropriately towards the BSA. Under extreme loading (only expected to occur under network fault conditions), lower priority data services and/or HSI traffic are compromised in order to protect video and voice traffic. Classification of HSI traffic based on source network address or IEEE 802.1p marking allows the QoS information to be propagated to upstream or downstream nodes by network elements. See Table 4 for the descriptions.

The BSR performs service distribution routing based on guarantees required to deliver the service and associated content, rather than on individual subscribers. The BSR only needs to classify content based on the required forwarding class for a given BSA to ensure that each service's traffic receives the appropriate treatment towards the BSA.

Figure 3 **Downstream QoS Enablement**

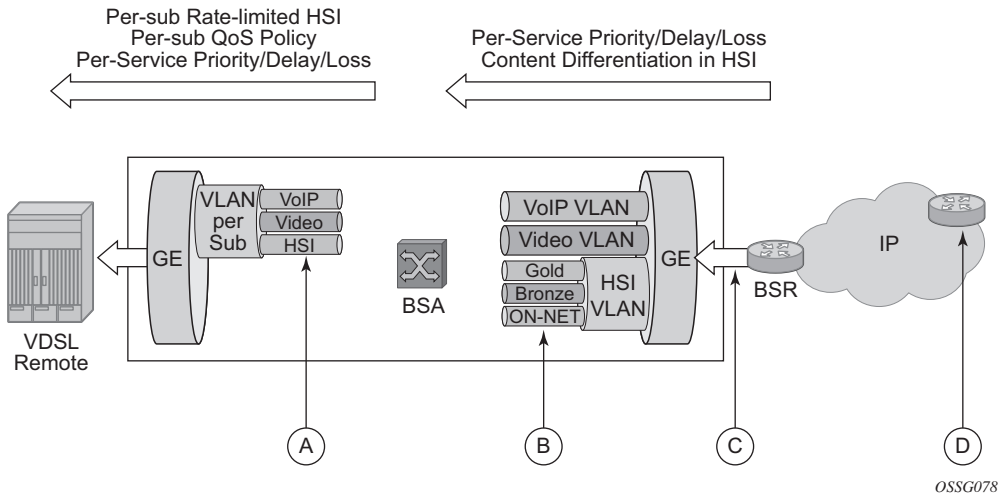

Table 4 **Downstream QoS Enablement**

Key	Description
A	Per-subscriber queuing and PIR/CIR policing/shaping for HSI. HSI service classified on source IP range. Per-service prioritization for VoIP and video. VoIP is prioritized over video. Destination IP and/or DSCP classification. 802.1 marking for prioritization in the access and home.
B	VoIP and video queued and prioritized on per-VLAN QoS policy basis. HSI content differentiation based on DSCP. Each queue may have individual CIR/PIR and shaping. Optical overall subscriber rate limiting on VLAN (H-QoS).
C	For HSI, content differentiation queuing for gold, silver, or bronze based on DSCP classification. Optional overall subscriber rate limiting on VLAN.
D	Preferred content marked (DSCP) of trusted ingress points of IP network.

In the upstream direction (BSA to BSR, as shown in Figure 4), traffic levels are substantially lower. Class-based queuing is used on the BSA network interface to ensure that video control traffic is propagated with a minimal and consistent delay, and that preferred data or HSI services receive better treatment for upstream/peering service traffic than the best effort Internet class of service.

Note: The IP edge is no longer burdened with enforcing per-subscriber policies for hundreds of thousands of users. This function is now distributed to the BSAs, and the per-subscriber policies can be implemented on the interfaces directly facing the access nodes.

Figure 4 **Upstream QoS Enablement**

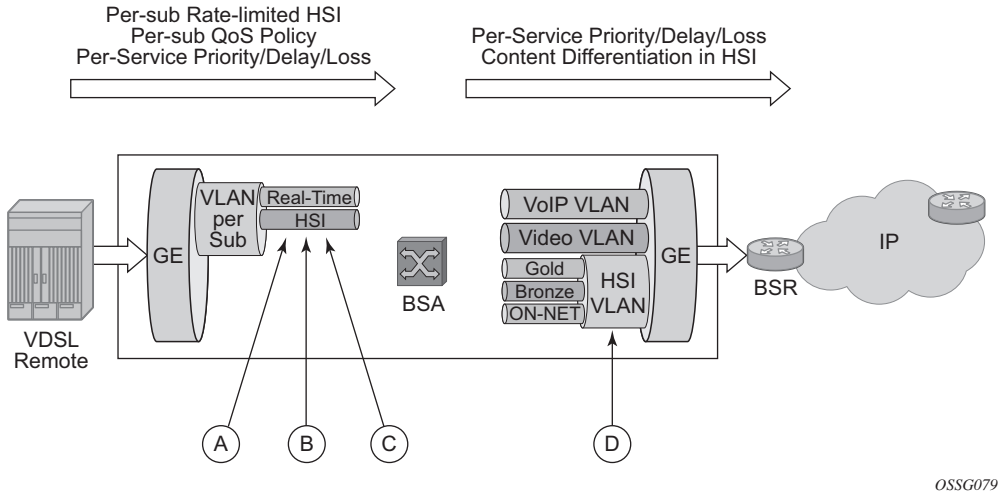

Upstream QoS enablement keys are described in Table 5.

Table 5 **Upstream QoS Enablement**

Key	Description
A	HSI: Per-subscriber queuing with PIR or CIR policy or shaping.
B	VoIP and Video: Shared queuing for prioritization of real-time traffic over HSI. Upstream video is negligible.
C	Per-subscriber QoS/Content classification for content differentiation.
D	Video and VoIP: Policy defines priority aggregate CIR and PIR. HSI: QoS policy defines priority and aggregate CIR and PIR. Content differentiation based on ingress classification. DSCP is marked.

The BSA is capable of scheduling and queuing functions on a per-service, per-subscriber basis, in addition to performing wire-speed packet classification and filtering based on both Layer 2 and Layer 3 fields.

Each subscriber interface provides at least three dedicated queues. TPSDA makes it possible to configure these queues such that the forwarding classes defined for all services can all be mapped to one service VLAN upstream. In the BSA, assuming hundreds of subscribers per gigabit Ethernet interface, this translates to a thousand or more queues per port.

In addition to per-service rate limiting for HSI services, each subscriber's service traffic can be rate limited as an aggregate using a bundled service policy. This allows different subscribers to receive different service levels independently and simultaneously.

Distributed multicasting today's predominant video service is broadcast TV, and will remain significant for a long time. As video services are introduced, it is sensible to optimize investments by matching resources to the service model relevant at the time. Consequently, the objective of the service infrastructure should be to incorporate sufficient flexibility to optimize for broadcast TV in the short term, yet scale to support a full unicast (VoD) model as video service offerings evolve.

Optimizing for broadcast TV means implementing multicast packet replication throughout the network. Multicast improves the efficiency of the network by reducing the bandwidth and fiber needed to deliver broadcast channels to the subscriber. A multicasting node can receive a single copy of a broadcast channel and replicate it to any downstream nodes that require it, substantially reducing the required network resources. This efficiency becomes increasingly important closer to the subscriber. Multicast should be performed at each or either of the access, aggregation, and video edge nodes.

Multicasting as close as possible to the subscriber has other benefits since it enables a large number of users to view the content concurrently. The challenges of video services are often encountered in the boundary cases, such as live sports events and breaking news, for which virtually all the subscribers may be watching just a few channels. These exceptional cases generally involve live content, which is true broadcast content. Multicasting throughout the network makes it possible to deliver content under these circumstances while simplifying the engineering of the network.

Efficient multicasting requires the distribution of functions throughout the access and the aggregation network to avoid overloading the network capacity with unnecessary traffic. TPSDA realizes efficient multicasting by implementing IGMP snooping in the access nodes, IGMP snooping in the BSA and multicast routing in the BSR (Figure 5).

Figure 5 **Distributed Multicasting in TPSDA**

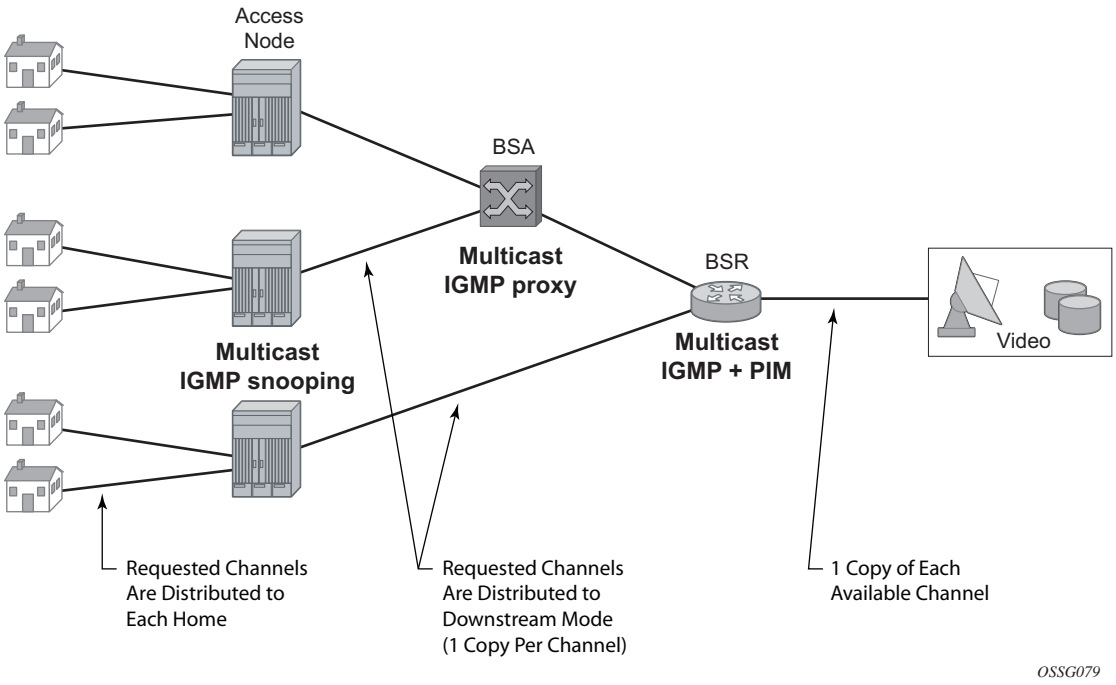

2.1.4.3 Virtual MAC Subnetting for VPLS

This feature allows, at the VPLS instance level, MAC subnetting, such as learning and switching based on a configurable number of bits from the source MAC address and from the destination MAC, respectively. This considerably reduces the VPLS FDB size.

MAC scalability involving MAC learning and switching based on the first x bits of a virtual MAC address is suitable in an environment where some MAC addresses can be aggregated based on a common first x bits, for example 28 out of 48. This can be deployed in a TPSDA environment where the VPLS is used for pure aggregation (there is no subscriber management) between the DSLAM and BRAS devices. The DSLAMs must be able to map customer MAC addresses to a pool of internal virtual MAC addresses where the first bits (28, for example) identify the DSLAM with the next 20 bits identifying, the DSLAM slot, port number, and customer MAC station on that port. The VPLS instance(s) in the PE distinguishes only between different DSLAMs connected to it. They need to learn and switch based only on the first 28 bits of the MAC address allowing scaling of the FDB size in the PE.

Figure 6 **Subnetting Topology**

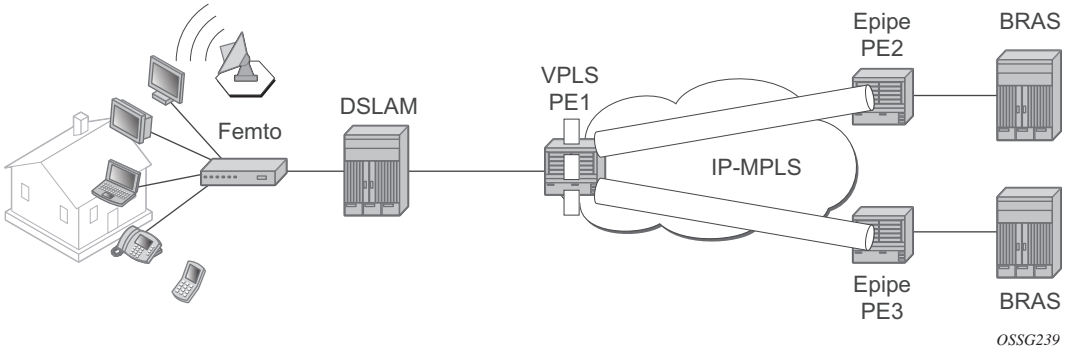

Figure 6 displays a Layer 2 PE network (such as the ESS-Series) aggregating traffic from DSLAMs (Nokia) to BRAS devices. The VPLS service is running in the PEs directly connected to the DSLAMs (VPLS PE1) while the PEs connected to the BRAS devices are running a point-to-point Layer 2 service (Epipe).

Nokia DSLAMs have the capability to map every customer MAC to a service provider MAC using the virtual MAC addressing scheme depicted in Figure 7.

Figure 7 **VMAC Subnetting Topology**

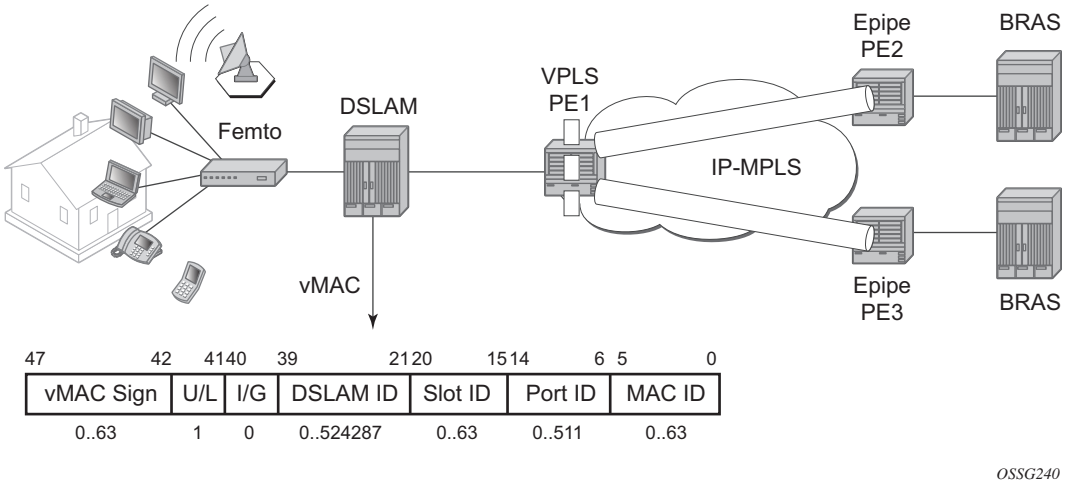

As the packet ingresses the DSLAM from the residential customer, the source MAC address (a customer MAC for one of its terminals or routers) is replaced by the DSLAM with a virtual MAC using the format depicted in Figure 7.

- The U/L bit is the seventh bit of the first byte and is used to determine whether the address is universally or locally administered. The U/L bit is set to 1 to indicate the address is locally administered.
- The following bits are used to build the VMAC address: DSLAM ID bits 39 — 21, slot ID bits 20 — 15, port ID bits 14-6 and the customer station ID bits 5 — 0.

Based on this scheme, it is apparent that the VMACs from one DSLAM have bits 47-21 in common.

The VPLS instance in PE1 only learns the first part of the MAC (bits 47 — 21) and, as the packets arrive from the BRAS device, switches based only on these bits in the destination MAC address to differentiate between the connected DSLAMs. Once the packet arrives at the DSLAM, the entire destination MAC is checked to determine the slot, port and which specific customer station the packet is destined to. As the packet is sent to the customer, the DSLAM replaces the destination MAC address with the actual customer MAC corresponding to the customer station.

The following are VPLS features not supported when the VMAC subnetting feature is enabled:

- Blocked features — CLI consistency checked provided
- Residential Split Horizon Groups
- BGP AD
- TPSDA (subscriber management) features

- PBB
- VPLS OAM (MAC populate, MAC ping, MAC trace, CPE Ping)

2.2 Services

A service is a globally unique entity that refers to a type of connectivity service for either Internet or VPN connectivity. Each service is uniquely identified by a service ID within a service area. The service model uses logical service entities to construct a service. In the service model, logical service entities are provide a uniform, service-centric configuration, management, and billing model for service provisioning.

Services can provide a Layer 2 or bridged service or a Layer 3 or IP routed connectivity between a service access point (SAP) on one router and another service access point (a SAP is where traffic enters and exits the service) on the same (local) or another 7450 ESS or 7750 SR OS (distributed). A distributed service spans more than one router.

Distributed services use service distribution points (SDPs) to direct traffic to another 7450 ESS or 7750 SR OS through a service tunnel. SDPs are created on each participating 7450 ESS or 7750 SR OS, specifying the origination address (the router participating in the service communication) and the destination address of another 7450 ESS or 7750 SR OS. SDPs are then bound to a specific customer service. Without the binding process, far-end 7450 ESS and 7750 SR devices are not able to participate in the service (there is no service without associating an SDP with a service).

2.2.1 Service Types

The 7450 ESS and 7750 SR OS offers the following types of subscriber services which are described in more detail in the referenced chapters:

- Virtual Leased Line (VLL) services:
  - Ethernet pipe (Epipe) — A Layer 2 point-to-point VLL service for Ethernet frames.
  - ATM VLL (Apipe) — A point-to-point ATM service between users connected to 7750 nodes on an IP or MPLS network.
  - Frame-Relay (Fpipe) — A point-to-point Frame Relay service between users connected to 7750 nodes on the IP or MPLS network.
  - IP Pipe (Ipipe) — Provides IP connectivity between a host attached to a point-to-point access circuit (FR, ATM, PPP) with routed IPv4 encapsulation and a host attached to an Ethernet interface.

- Virtual Private LAN Service (VPLS) — A Layer 2 multipoint-to-multipoint VPN. VPLS includes Hierarchical VPLS (H-VPLS) which is an enhancement of VPLS which extends Martini-style signaled or static virtual circuit labeling outside the fully meshed VPLS core.
- Internet Enhanced Service (IES) — A direct Internet access service where the customer is assigned an IP interface for Internet connectivity.
- Virtual Private Routed Network (VPRN) — Layer 3 IP multipoint-to-multipoint VPN service as defined in RFC 2547bis. VPRN is supported on the 7750 SR only.

2.2.2 Service Policies

Common to all 7450 ESS and 7750 SR connectivity services are policies that are assigned to the service. Policies are defined at a global level and then applied to a service on the router. Policies are used to define service enhancements. The types of policies that are common to all 7450 ESS and 7750 SR connectivity services are:

- SAP Quality of Service (QoS) policies which allow for different classes of traffic within a service at SAP ingress and SAP egress.

  QoS ingress and egress policies determine the QoS characteristics for a SAP. A QoS policy applied to a SAP specifies the number of queues, queue characteristics (such as forwarding class, committed, and peak information rates, and so on) and the mapping of traffic to a forwarding class. A QoS policy must be created before it can be applied to a SAP. A single ingress and a single egress QoS policy can be associated with a SAP.
- Filter policies allow selective blocking of traffic matching criteria from ingressing or egressing a SAP.

  Filter policies, also referred to as access control lists (ACLs), control the traffic allowed in or out of a SAP based on MAC or IP match criteria. Associating a filter policy on a SAP is optional. Filter policies are identified by a unique filter policy ID. A filter policy must be created before it can be applied to a SAP. A single ingress and single egress filter policy can be associated with a SAP.
- Scheduler policies define the hierarchy and operating parameters for virtual schedulers. Schedulers are divided into groups based on the tier each scheduler is created under. A tier is used to give structure to the schedulers within a policy and define rules for parent scheduler associations.
- Accounting policies define how to count the traffic usage for a service for billing purposes.

The 7450 ESS and 7750 SR OS provide a comprehensive set of service-related counters. Accounting data can be collected on a per-service, per-forwarding class basis, which enables network operators to accurately measure network usage and bill each customer for each individual service using any of a number of different billing models.

2.3 Nokia Service Model

2.3.1 Introduction

In the 7450 ESS and 7750 SR service model, the service edge routers are deployed at the provider edge. Services are provisioned on the routers and transported across an IP and/or IP/MPLS provider core network in encapsulation tunnels created using Generic Router Encapsulation (GRE) or MPLS Label Switched Paths (LSPs).

The service model uses logical service entities to construct a service. The logical service entities are designed to provide a uniform, service-centric configuration, management, and billing model for service provisioning. Some benefits of this service-centric design include:

- Many services can be bound to a single customer.
- Many services can be bound to a single tunnel.
- Tunnel configurations are independent of the services they carry.
- Changes are made to a single logical entity rather than multiple ports on multiple devices. It is easier to change one tunnel rather than several services.
- The operational integrity of a logical entity (such as a service tunnel and service end points) can be verified rather than dozens of individual services improving management scaling and performance.
- A failure in the network core can be correlated to specific subscribers and services.
- QoS policies, filter policies, and accounting policies are applied to each service instead of correlating parameters and statistics from ports to customers to services.

Service provisioning uses logical entities to provision a service where additional properties can be configured for bandwidth provisioning, QoS, security filtering, accounting or billing to the appropriate entity.

2.3.2 Service Entities

The basic logical entities in the service model used to construct a service are:

- Customers
- Service Access Points (SAPs)

• Service Distribution Points (SDPs) (for distributed services only)

Nokia's service entities are shown in Figure 8.

Figure 8 **Nokia's Service Entities**

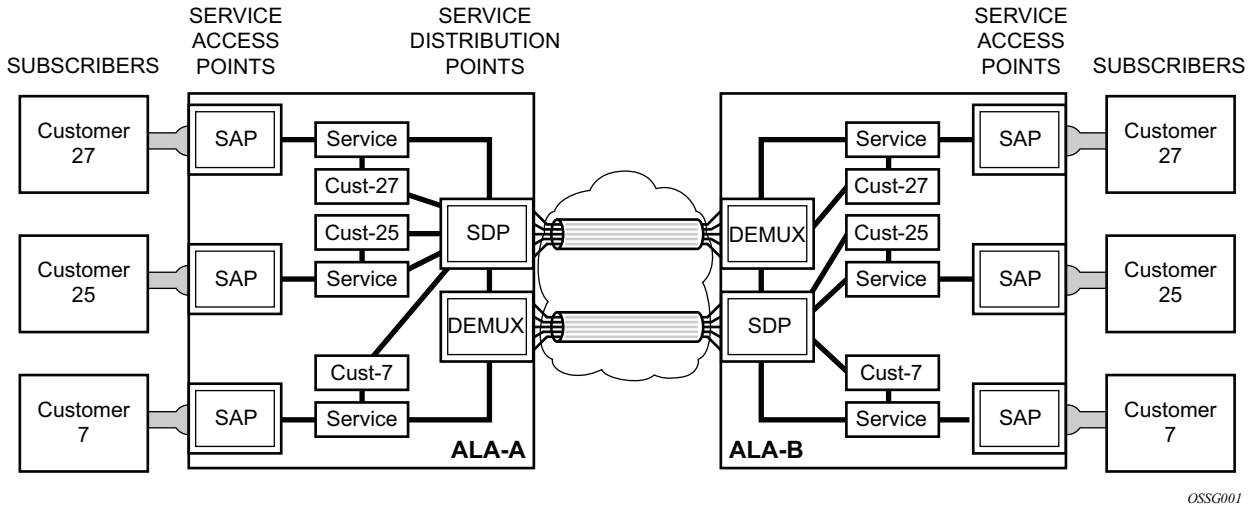

2.3.3 Customers

The terms customers and subscribers are used synonymously. The most basic required entity is the customer ID value which is assigned when the customer account is created. To provision a service, a customer ID must be associated with the service at the time of service creation.

2.3.4 Service Access Points (SAPs)

Each subscriber service type is configured with at least one service access point (SAP). A SAP identifies the customer interface point for a service on an Nokia 7450 ESS or 7750 SR OS (Figure 9). The SAP configuration requires that slot, MDA, and port/channel (7750 SR) information be specified. The slot, MDA, and port/channel parameters must be configured prior to provisioning a service (refer to the Cards, MDAs, and Ports section of the *7450 ESS, 7750 SR, 7950 XRS, and VSR Interface Configuration Guide*).

A SAP is a local entity to the 7450 ESS and 7750 SR and is uniquely identified by:

- The physical Ethernet port or SONET/SDH port or TDM channel (7750 SR)
- The encapsulation type
- The encapsulation identifier (ID)

Depending on the encapsulation, a physical port or channel can have more than one SAP associated with it. SAPs can only be created on ports or channels designated as “access” in the physical port configuration. SAPs cannot be created on ports designated as core-facing “network” ports as these ports have a different set of features enabled in software.

Figure 9 **Service Access Point (SAP)**

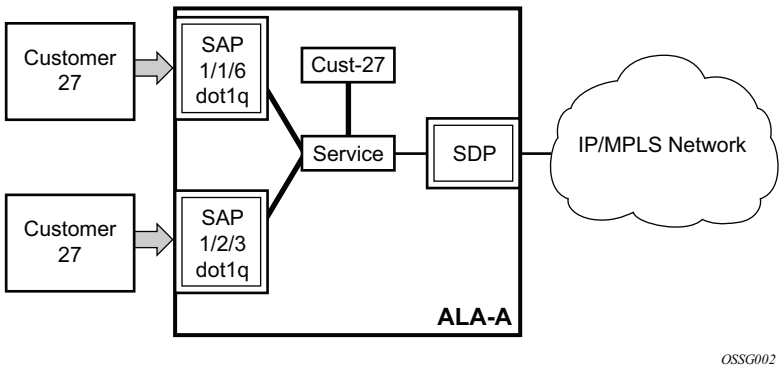

2.3.4.1 SAP Encapsulation Types and Identifiers

The encapsulation type is an access property of a service Ethernet port or SONET/SDH or TDM channel (the TDM channel applies to the 7750 SR). The appropriate encapsulation type for the port or channel depends on the requirements to support multiple services on a single port/channel on the associated SAP and the capabilities of the downstream equipment connected to the port or channel. For example, a port can be tagged with IEEE 802.1Q (referred to as dot1q) encapsulation in which each individual tag can be identified with a service. A SAP is created on a given port or channel by identifying the service with a specific encapsulation ID.

2.3.4.2 Ethernet Encapsulations

The following lists encapsulation service options on Ethernet ports:

1. Null — Supports a single service on the port. For example, where a single customer with a single service customer edge (CE) device is attached to the port. The encapsulation ID is always 0 (zero).

2. Dot1q — Supports multiple services for one customer or services for multiple customers (Figure 10). For example, the port is connected to a multi-tenant unit (MTU) device with multiple downstream customers. The encapsulation ID used to distinguish an individual service is the VLAN ID in the IEEE 802.1Q header.
3. Q-in-Q — The q-in-q encapsulation type adds a IEEE 802.1Q tag to the 802.1Q tagged packets entering the network to expand the VLAN space by tagging tagged packets, producing a double tagged frame.

Note: The SAP can be defined with a wildcard for the inner label. (for example, “100:*”). In this situation all packets with an outer label of 100 are treated as belonging to the SAP. If, on the same physical link there is also a SAP defined with a q-in-q encap of 100:1, then traffic with the 100:1 encap will go to that SAP and all other traffic with 100 as the first label will go to the SAP with the 100:* definition.

In the dot1q and q-in-q options, traffic encapsulated with tags for which there is no definition are discarded.

Figure 10 **Multiple SAPs on a Single Port or Channel**

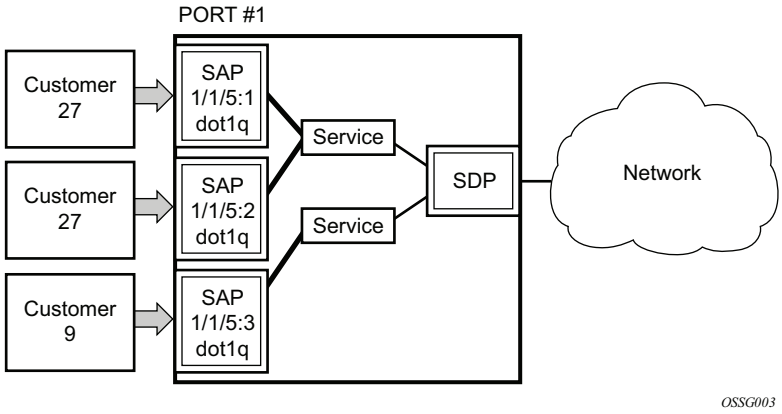

2.3.4.3 SAP Considerations

When configuring a SAP, consider the following:

- A SAP is a local entity and only locally unique to a given device. The same SAP ID value can be used on another 7450 ESS or 7750 SR.
- There are no default SAPs. All SAPs must be created.
- The default administrative state for a SAP at creation time is administratively enabled.

- When a SAP is deleted, all configuration parameters for the SAP is also deleted. For Internet Ethernet Service (IES), the IP interface must be shutdown before the SAP on that interface may be removed.
- A SAP is owned by and associated with the service in which it is created in each 7450 ESS or 7750 SR.
- A port or channel with a dot1q or BCP-dot1q encapsulation type means the traffic for the SAP is identified based on a specific IEEE 802.1Q VLAN ID value. The VLAN ID is stripped off at SAP ingress and the appropriate VLAN ID placed on at SAP egress. As a result, VLAN IDs only have local significance, so the VLAN IDs for the SAPs for a service need not be the same at each SAP.
- If a port or channel is administratively shutdown, all SAPs on that port or channel are operationally out of service.
- A SAP cannot be deleted until it has been administratively disabled (shutdown).
- Each SAP can be configured with only the following:
  - Ingress or egress filter policy
  - Ingress or egress QoS policy
  - Accounting policy
  - Ingress or egress scheduler policy

2.3.5 Service Distribution Points (SDPs)

A service distribution point (SDP) acts as a logical way to direct traffic from one 7450 ESS or 7750 SR to another 7450 ESS or 7750 SR through a uni-directional (one-way) service tunnel. The SDP terminates at the far-end router which directs packets to the correct service egress SAPs on that device. A distributed service consists of a configuration with at least one SAP on a local node, one SAP on a remote node, and an SDP binding the service to the service tunnel.

An SDP has the following characteristics:

- An SDP is locally unique to a participating 7450 ESS or 7750 SR. The same SDP ID can appear on other 7450 ESS or 7750 SR OS.
- An SDP uses the system IP address to identify the far-end 7450 ESS or 7750 SR edge router.
- An SDP is not specific to any one service or any type of service. Once an SDP is created, services are bound to the SDP. An SDP can also have more than one service type associated with it.
- All services mapped to an SDP use the same transport encapsulation type defined for the SDP (either GRE or MPLS).

- An SDP is a management entity. Even though the SDP configuration and the services carried within are independent, they are related objects. Operations on the SDP affect all the services associated with the SDP. For example, the operational and administrative state of an SDP controls the state of services bound to the SDP.

An SDP from the local device to a far-end 7450 ESS or 7750 SR requires a return path SDP from the far-end router back to the local router. Each device must have an SDP defined for every remote router to which it wants to provide service. SDPs must be created first, before a distributed service can be configured.

2.3.5.1 SDP Binding

To configure a distributed service from ALA-A to ALA-B, the SDP ID (4) (Figure 11) must be specified in the service creation process in order to “bind” the service to the tunnel (the SDP). Otherwise, service traffic is not directed to a far-end point and the far-end 7450 ESS and 7750 SR device(s) cannot participate in the service (there is no service).

Figure 11 **A GRE Service Distribution Point (SDP) pointing from ALA-A to ALA-B**

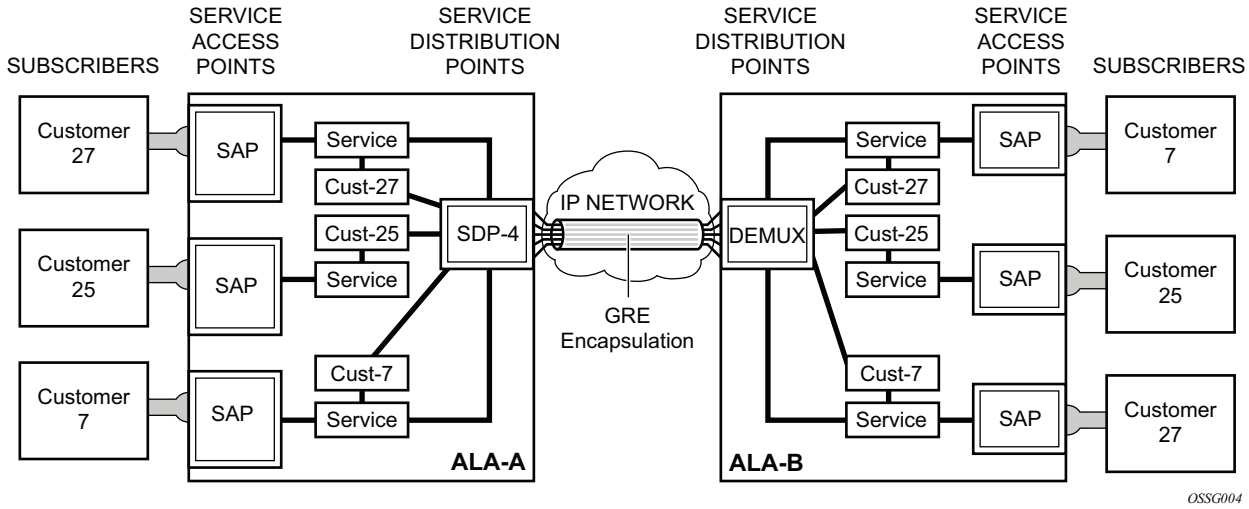

2.3.5.2 Spoke and Mesh SDPs

When an SDP is bound to a service, it is bound as either a spoke SDP or a mesh SDP. The type of SDP indicates how flooded traffic is transmitted.

A spoke SDP is treated like the equivalent of a traditional bridge “port” where flooded traffic received on the spoke SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

All mesh SDPs bound to a service are logically treated like a single bridge “port” for flooded traffic where flooded traffic received on any mesh SDP on the service is replicated to other “ports” (spoke SDPs and SAPs) and not transmitted on any mesh SDPs.

2.3.5.3 SDP Encapsulation Types

The Nokia service model uses encapsulation tunnels through the core to interconnect 7450 ESS and 7750 SR service edge routers. An SDP is a logical way of referencing the entrance to an encapsulation tunnel.

The following encapsulation types are supported:

- Layer 2 within Generic Routing Encapsulation (GRE)
- Layer 2 within RSVP signaled, loose hop non-reserved MPLS LSP
- Layer 2 within RSVP signaled, strict hop non-reserved MPLS LSP
- Layer 2 within RSVP-TE signaled, bandwidth reserved MPLS LSP

2.3.5.3.1 GRE

GRE encapsulated tunnels have very low overhead and are best used for Best-Effort class of service. Packets within the GRE tunnel follow the Interior Gateway Protocol (IGP) shortest path from edge to edge. If a failure occurs within the service core network, the tunnel only converges as quickly as the IGP itself. If Equal Cost Multi-Path (ECMP) routing is used in the core, many loss-of-service failures can be minimized to sub-second timeframes.

2.3.5.3.2 MPLS

Multi-Protocol Label Switching (MPLS) encapsulation has the following characteristics:

- LSPs (label switched paths) are used through the network, for example, primary, secondary, loose hop, and so on These paths define how traffic traverses the network from point A to B. If a path is down, depending on the configuration parameters, another path is substituted.

Paths can be manually defined or a constraint-based routing protocol (such as OSPF-TE or CSPF) can be used to determine the best path with specific constraints.

- The 7450 ESS and 7750 SR OS support both signaled and non-signaled LSPs through the network.
- Non-signaled paths are defined at each hop through the network.
- Signaled paths are communicated via protocol from end to end using Resource Reservation Protocol (RSVP).

  Because services are carried in encapsulation tunnels and an SDP is an entrance to the tunnel, an SDP has an implicit Maximum Transmission Unit (MTU) value. The MTU for the service tunnel can affect and interact with the MTU supported on the physical port where the SAP is defined.

2.3.5.4 SDP Keepalives

SDP keepalives are a way of actively monitoring the SDP operational state using periodic Nokia SDP Ping Echo Request and Echo Reply messages. Nokia SDP Ping is a part of Nokia's suite of Service Diagnostics built on an Nokia service-level OA&M protocol. When SDP Ping is used in the SDP keepalive application, the SDP Echo Request and Echo Reply messages are a mechanism for exchanging far-end SDP status.

Configuring SDP keepalives on a given SDP is optional. SDP keepalives for a particular SDP have the following configurable parameters:

- AdminUp or AdminDown State
- Hello Time
- Message Length
- Max Drop Count
- Hold Down Time

SDP keepalive Echo Request messages are only sent when the SDP is completely configured and administratively up and SDP keepalives is administratively up. If the SDP is administratively down, keepalives for the SDP are disabled.

SDP keepalive Echo Request messages are sent out periodically based on the configured Hello Time. An optional Message Length for the Echo Request can be configured. If Max Drop Count Echo Request messages do not receive an Echo Reply, the SDP is immediately brought operationally down.

If a keepalive response is received that indicates an error condition, the SDP is immediately brought operationally down.

Once a response is received that indicates the error has cleared and the Hold Down Time interval has expired, the SDP is eligible to be put into the operationally up state. If no other condition prevents the operational change, the SDP enters the operational state.

2.4 Epipe Service Overview

An Epipe service is Nokia's implementations of an Ethernet VLL based on the IETF Martini Drafts (*draft-martini-l2circuit-trans-mpls-08.txt* and *draft-martini-l2circuit-encapmpls-04.txt*) and the IETF Ethernet Pseudo-wire Draft (*draft-so-pwe3-ethernet-00.txt*).

An Epipe service is a Layer 2 point-to-point service where the customer data is encapsulated and transported across a service provider's IP or MPLS network. An Epipe service is completely transparent to the subscriber's data and protocols. The 7450 ESS and 7750 SR Epipe service does not perform any MAC learning. A local Epipe service consists of two SAPs on the same node, whereas a distributed Epipe service consists of two SAPs on different nodes. SDPs are not used in local Epipe services.

Each SAP configuration includes a specific port or channel (applies to the 7750 SR) on which service traffic enters the router from the customer side (also called the access side). Each port is configured with an encapsulation type. If a port is configured with an IEEE 802.1Q (referred to as dot1q) encapsulation, then a unique encapsulation value (ID) must be specified.

The Epipe/VLL service is shown in Figure 12.

Figure 12 **Epipe/VLL Service**

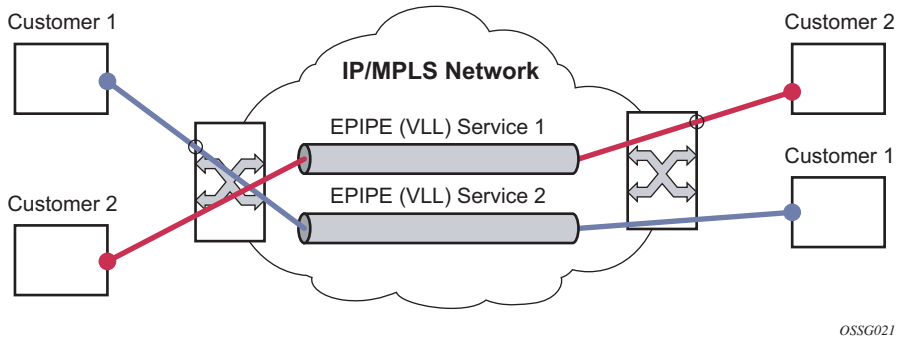

2.5 VPLS Service Overview

Virtual Private LAN Service (VPLS) as described in Internet Draft *draft-ietf-ppvpn-vpls-ldp-01.txt*, is a class of virtual private network service that allows the connection of multiple sites in a single bridged domain over a provider-managed IP/MPLS network. The customer sites in a VPLS instance appear to be on the same LAN, regardless of their location. VPLS uses an Ethernet interface on the customer-facing (access) side which simplifies the LAN/WAN boundary and allows for rapid and flexible service provisioning.

VPLS offers a balance between point-to-point Frame Relay service and outsourced routed services (VPRN). VPLS enables each customer to maintain control of their own routing strategies. All customer routers in the VPLS service are part of the same subnet (LAN) which simplifies the IP addressing plan, especially when compared to a mesh constructed from many separate point-to-point connections. The VPLS service management is simplified since the service is not aware of nor participates in the IP addressing and routing.

A VPLS service provides connectivity between two or more SAPs on one (which is considered a local service) or more (which is considered a distributed service) routers. The connection appears to be a bridged domain to the customer sites so protocols, including routing protocols, can traverse the VPLS service.

Other VPLS advantages include:

- VPLS is a transparent, protocol-independent service.
- There is no Layer 2 protocol conversion between LAN and WAN technologies.
- There is no need to design, manage, configure, and maintain separate WAN access equipment, thus, eliminating the need to train personnel on WAN technologies such as Frame Relay.

For details on VPLS, including a packet walkthrough, refer to VPLS section in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN*.

2.5.1 Split Horizon SAP Groups and Split Horizon Spoke SDP Groups

Within the context of VPLS services, a loop-free topology within a fully meshed VPLS core is achieved by applying split-horizon forwarding concept that packets received from a mesh SDP are never forwarded to other mesh SDPs within the same service. The advantage of this approach is that no protocol is required to detect loops within the VPLS core network.

In applications such as DSL aggregation, it is useful to extend this split-horizon concept also to groups of SAPs and/or spoke SDPs. This extension is referred to as a split horizon SAP group or residential bridging.

Traffic arriving on a SAP or a spoke SDP within a split horizon group is not copied to other SAPs and spoke SDPs in the same split horizon group (but is copied to SAPs or spoke SDPs in other split horizon groups if these exist within the same VPLS).

2.5.1.1 Residential Split Horizon Groups

To improve the scalability of a SAP-per-subscriber model in the broadband services aggregator (BSA), the 7450 ESS and 7750 SR support a variant of split horizon groups called residential split horizon groups (RSHG).

A RSHG is a group of split horizon group SAPs with following limitations:

- Downstream broadcast traffic is not allowed.
- Downstream multicast traffic is allowed when IGMP snooping is configured in the VPLS.
- STP is not supported.

Spoke SDPs can also be members of a RSHG VPLS. The downstream multicast traffic restriction does not apply to spoke SDPs.

2.6 IES Service Overview

Internet Enhanced Service (IES) is a routed connectivity service where the subscriber communicates with an IP router interface to send and receive Internet traffic. An IES has one or more logical IP routing interfaces each with a SAP which acts as the access point to the subscriber's network. IES allow customer-facing IP interfaces to participate in the same routing instance used for service network core routing connectivity. IES services require that the IP addressing scheme used by the subscriber be unique between other provider addressing schemes and potentially the entire Internet.

While IES is part of the routing domain, the usable IP address space may be limited. This allows a portion of the service provider address space to be reserved for service IP provisioning, and be administered by a separate but subordinate address authority.

IP interfaces defined within the context of an IES service must have a SAP associated as the access point to the subscriber network, as shown in Figure 13. Multiple IES services are created to segregate subscriber-owned IP interfaces.

Figure 13 **Internet Enhanced Service**

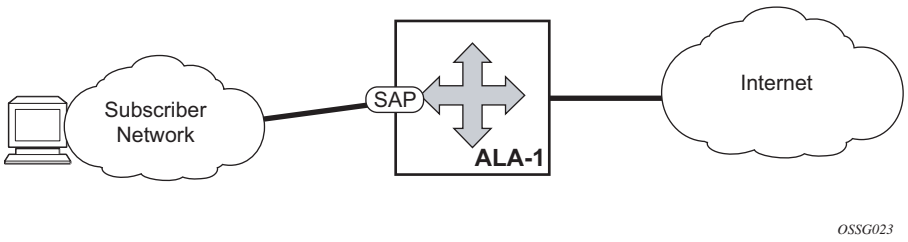

The IES service provides Internet connectivity. Other features include:

- Multiple IES services are created to separate customer-owned IP interfaces.
- More than one IES service can be created for a single customer ID.
- More than one IP interface can be created within a single IES service ID. All IP interfaces created within an IES service ID belong to the same customer.

2.6.1 IP Interface

IES customer IP interfaces can be configured with most of the same options found on the core IP interfaces. The advanced configuration options supported are:

- VRRP (for IES services with more than one IP interface)
- Cflowd (supported on the 7750 SR only)
- Secondary IP addresses
- ICMP options

Configuration options found on core IP interfaces not supported on IES IP interfaces are:

- Unnumbered interfaces
- NTP broadcast receipt

2.7 VPRN Service Overview

VPRN service is supported on the 7750 SR only.

RFC2547bis is an extension to the original RFC 2547, which details a method of distributing routing information and forwarding data to provide a Layer 3 Virtual Private Network (VPN) service to end customers.

Each Virtual Private Routed Network (VPRN) consists of a set of customer sites connected to one or more PE routers. Each associated PE router maintains a separate IP forwarding table for each VPRN. Additionally, the PE routers exchange the routing information configured or learned from all customer sites via MP-BGP peering. Each route exchanged via the MP-BGP protocol includes a Route Distinguisher (RD), which identifies the VPRN association.

The service provider uses BGP to exchange the routes of a particular VPN among the PE routers that are attached to that VPN. This is done in a way which ensures that routes from different VPNs remain distinct and separate, even if two VPNs have an overlapping address space. The PE routers distribute routes from other CE routers in that VPN to the CE routers in a particular VPN. Since the CE routers do not peer with each other there is no overlay visible to the VPN's routing algorithm.

When BGP distributes a VPN route, it also distributes an MPLS label for that route. On a SR-Series, a single label is assigned to all routes in a VPN.

Before a customer data packet travels across the service provider's backbone, it is encapsulated with the MPLS label that corresponds, in the customer's VPN, to the route which best matches the packet's destination address. The MPLS packet is further encapsulated with either another MPLS label or GRE tunnel header, so that it gets tunneled across the backbone to the proper PE router. Each route exchanged by the MP-BGP protocol includes a route distinguisher (RD), which identifies the VPRN association. Thus the backbone core routers do not need to know the VPN routes.

2.8 Deploying and Provisioning Services

The service model provides a logical and uniform way of constructing connectivity services. The basic steps for deploying and provisioning services can be broken down into three phases.

2.8.1 Phase 1: Core Network Construction

Before the services are provisioned, the following tasks should be completed:

- Build the IP or IP/MPLS core network.
- Configure routing protocols.
- Configure MPLS LSPs (if MPLS is used).
- Construct the core SDP service tunnel mesh for the services.

2.8.2 Phase 2: Service Administration

Perform preliminary policy and SDP configurations to control traffic flow, operator access, and to manage fault conditions and alarm messages, the following tasks should be completed:

- Configure group and user access privileges.
- Build templates for QoS, filter and accounting policies needed to support the core services.

2.8.3 Phase 3: Service Provisioning

- Provision customer account information.
- If necessary, build any customer-specific QoS, filter or accounting policies.
- Provision the customer services on the 7450 ESS or 7750 SR service edge routers by defining SAPs, binding policies to the SAPs, and then binding the service to appropriate SDPs as necessary.

2.9 Configuration Notes

This section describes service configuration restrictions.

2.9.1 General

Service provisioning tasks can be logically separated into two main functional areas, core tasks and subscriber tasks and are typically performed prior to provisioning a subscriber service.

Core tasks include the following:

- Create customer accounts
- Create template QoS, filter, scheduler, and accounting policies
- Create LSPs
- Create SDPs

Subscriber services tasks include the following:

- Create VLL, VPLS, IES, or VPRN services (VPRN is supported on the 7750 SR only)
- Configure interfaces (where required) and SAPs
- Bind SDPs
- Create exclusive QoS and filter policies

2.10 Configuring Triple Play Services with CLI

This section provides information to configure Residential Broadband Aggregation services using the command line interface. It is assumed that the reader is familiar with basic configuration of VPLS, IES and VPRN services (VPRN is supported on the 7750 SR only).

2.10.1 Configuring VPLS Residential Split Horizon Groups

To configure a group of SAPs in a VPLS service as a Residential Split Horizon Group (RSHG), add the residential-group parameter when creating the split horizon group. Traffic arriving on a SAP within an RSHG is not copied to other SAPs in the same split horizon group.

Note: The split horizon group must be created before it can be applied.

The following example displays a VPLS configuration with split horizon enabled:

```
*A:ALA-48>config>service>vpls# info
----------------------------------------------
            split-horizon-group "DSL-group2" residential-group create
                description "split horizon group for DSL - no broadcast supported"
            exit
            stp
                shutdown
            exit
            sap 2/1/4:100 split-horizon-group "DSL-group2" create
                description "SAP in RSHG"
            exit
            sap 2/1/4:200 split-horizon-group "DSL-group2" create
                description "another SAP in the RSHG"
            exit
            no shutdown
----------------------------------------------
*A:ALA-48>config>service>vpls#
```

2.10.1.1 Configuring Static Hosts

In order for the static host to be operational, forwarding traffic bi-directional, the mac address of the host must be learned or configured. Learning the MAC of the static host is different for IPv4 or IPv6.

If an IPv4 static host MAC is not specified

- The system learns respective MAC address dynamically from ARP packets (arp-request or gratuitous-arp) generated by the host with the specified IP address.
- On a VPLS service, this can occur if arp-reply-agent function is enabled on a given SAP. On Layer 3 services, such as IES or VPRN, the ARP packets are always examined so no further conditions are applicable.

If an IPv6 static host MAC is not specified

- The system learns the MAC address depending on the type of host configured such as: IPv6 prefix host or IPv6 address host.
  - A SAP can be specified as a single-MAC and it implies that there is only a single device attached to the SAP. It changes the MAC learning behavior on the SAP for IPv6 host only. Firstly, all IPv6 hosts shares the same learned MAC. Secondly, the MAC address is learned from the host's router solicits and neighbor discoveries.
  - For static host with an address, upon shutdown, a RS for the clients IPv6 address is sent towards the host and the MAC is learned upon the RA reply
  - For static host with an address, SHCV sends the RS, and the MAC is learned from the RA
  - For static host with an address, the OAM command can trigger a RS and the MAC is learned from the RA
  - For static host with either a prefix or an address, linking the IPv6 host to an IPv4 host copies the IPv4 host MAC address to the IPv6 host and vice versa.

The learned MAC address is handled as a MAC address of static host with explicitly defined mac-address. Meaning:

- The MAC address is not aged by the mac-aging or any other aging timers.
- The MAC address is not moved to a new SAP as a consequence of relearning event. A relearning event is a learning request for the same MAC address which comes from a new SAP.
- The MAC address is not flushed from FDB due to SAP failure or STP flush messages.

Every time the given static-host uses a different MAC address in its ARP request, the dynamic MAC learning process is performed. The old MAC address is overwritten by a new MAC address.

The learned MAC address is not made persistent (a static host is not a part of the persistency file). A service discontinuity of such a host could be proportional to its arp-cache timeout.

The following interactions are described:

- Antispoof (all services) — If a static IP-only host is configured on a given SAP, anti-spoof types, IP, NH MAC, and IP MAC are supported. Static hosts for which the MAC address is not known does not have an anti-spoof entry. It is added only after the corresponding MAC has been learned. As a consequence, all traffic generated by the host before the MAC is learned are dropped.
- MAC-linking (IES and VPRN service only) — The MAC address can be learned from either the IPv4 or IPv6 host. Once learned it is copied over to the host of the other address family
- Single-MAC (IES and VPRN service only) — This specifies that there is only one single subscriber (MAC) on the SAP and any ICMP6 message from the SAP can be assumed to be the subscriber MAC address. This does not apply to IPv4 host.
- Enhanced subscriber management (all services) — ESM is supported in a combination with a static ip-only host. It is assumed that ip-mac antispoofing is enabled. The resources (queues, and so on) are allocated at the time such a host is configured, although they are effectively used only after anti-spoof entry has been installed.
- Dual-homing (for IPv4 host only) — It is assumed that static host is configured on both chassis. The dynamic mac-address learning event is then synchronized (also, if the members are on two different nodes) and corresponding anti-spoof entries are installed on both chassis.
- MAC-pinning (for VPLS services only) — The dynamically learned MAC address of the static-host is considered as a static-mac and is not affected by the **no mac-pinning** command.
- ARP-reply-agent (VPLS services only) — It is possible to the enable arp-reply-agent on a SAP where static host with **ip-only** configured. In addition to the regular arp-reply-agent functionality (the reply to all arp-requests targeting the given host's IP address) learning of the host's MAC address is performed. As long as no MAC address have been learned no ARP replies on behalf of such host should be expected. Enabling of arp-reply-agent is optional for SAP with **ip-only** static hosts.

2.10.1.1.1 BNG Learning IP-Only Static Host’s MAC Address

Before an IP-only static host can forward or receive traffic, the MAC of the host must be learned by the BNG. The learned MAC is used to fill in the anti-spoof entry for the static host. Table 6 summarizes the different methods that the BNG could use to learn the IP only static host MAC address.

Table 6 **Methods to Learn IP-Only Static Host MAC Addresses**

Number of Host(s) on SAP	IPv4 Static Host	IPv6 Numbered Static Host	IPv6 Prefix Only Host (unnumbered)	Dual Stack Host
1	Host's ARP Triggered SHCV Data packet OAM SHCV	Auto populate RS and NS MAC GUA based NS Triggered SHCV Data packet OAM SHCV	Auto populate RS and NS MAC GUA based NS Data packet	IPv4: Via IPv6 static host MAC (if available) and mechanisms listed for IPv4 static host IPv6 numbered host: Via IPv4 static host MAC (if available) and mechanisms listed for IPv6 numbered static host IPv6 prefix only host: Via IPv4 static host MAC (if available) and mechanism listed for IPv6 prefix only host
N	Host's ARP Triggered SHCV Data packet OAM SHCV	GUA based NS SHCV Data packet OAM SHCV	GUA based NS Data packet	IPv4 Via IPv6 static host MAC (if available) and mechanism listed for IPv4 static host IPv6 numbered host: Via IPv4 static host MAC (if available) and mechanism listed for IPv6 numbered static host IPv6 prefix only host: Via IPv4 static host MAC (if available) and mechanism listed for IPv6 prefix only host

IP-only static hosts without a MAC address are put into a "non-forwarding" state and are not be able to send or receive traffic. The BNG must populate the MAC address using the methods described in the table above to put the host into a "forwarding" state. If the IPv4 static host sends an ARP, then the IPv4 address along with Ethernet header source MAC can populate the IP-only host MAC address. For IPv6, the subscriber must send in a NS using the GUA to allow the BNG to resolve the MAC address. If the IP-only host is an IPv6 prefix, NS using the GUA can also resolve the static host MAC address, provided that the GUA is within the subnet of the prefix. n IPv6 prefix host can only have one MAC association, multiple bridged hosts using a single IPv6 prefix is not supported. To support multiple bridge hosts on a SAP, individual /128 host must be configured along with enabling **ipoe-bridge-mode**.

2.10.1.1.2 Static Host Learning the IPv6 Default Gateway Address

Static hosts also need to resolve the default gateway and MAC in order to forward traffic. Static hosts can learn the BNG MAC using the following methods.

1. If the static host can use the BNG link-local as the default gateway, then the exchange of router solicit and advertisement should be able to populate the default route for the static host.
2. If the static host can use the BNG link-local as the default gateway but only through neighbor solicit, the BNG responds to NS only if the source address is the GUA of the host.
3. If the static host can only use a GUA as the default gateway, then the subscriber interface must be configured with an IPv6 address instead of a prefix.

A host's RS or NS messages might use link-local as the source address. The host's link-local addresses are unknown to the BNG, since the BNG only knows static hosts by their GUA. By default, unknown host's RS or NS messages are dropped. Therefore, static hosts cannot resolve their default gateway or MAC address. There are two different methods to enable the BNG to respond to static host's RS or NS messages. The first solution is to declare the SAP as a single-mac. When a SAP is configured as a "single-mac", it is declaring that only a single host is expected on the SAP and any RS or NS sent by the host automatically populates the MAC table. The BNG in return also responds to the static host's RS or NS. The second solution is for cases where multiple hosts are on the SAP. The RS or NS might not provide enough information to populate the subscriber host's MAC because the source IP are link-local addresses. The BNG anti-spoof only replies to known hosts RS or NS; therefore, these static hosts' RS or NS are dropped. The group interface has a configuration option to auto-reply to any RS or NS, which allows the static host to resolve their default gateway or route. Afterwards, the static host can send traffic to the BNG, which the BNG can use for MAC learning. To use data packets for MAC learning, the BNG has a "data-trigger" learning option for static hosts.

2.10.1.1.3 Configuring Static Hosts on an VPLS SAP

The following example displays a static host on a VPLS SAP configuration:

```
*A:ALA-48>config>service# info
----------------------------------------------
vpls 800 customer 6001 create
    description "VPLS with residential split horizon for DSL"
    stp
        shutdown
    exit
    sap 1/2/7:100 split-horizon-group "DSL-group2" create
        description "SAP for RSHG"
        static-host ip 10.1.1.1
```

```
    exit
    no shutdown
-----------------------------------------------
*A:ALA-48>config>service#
```

2.10.1.1.4 Configuring Static Hosts on an IES SAP

The following displays a static host on an IES SAP:

```
*A:ALA-49>config>service>ies>sub-if>grp-if# info
-----------------------------------------------
    sap 7/1/5 create
    description "IES with static host"
        static-host ip 10.1.1.1 create
        static-host ip 2001::1/128 create
        static-host ip 2001:1::/64 create
    exit
-----------------------------------------------
*A:ALA-49>config>service>ies#
```

2.10.1.1.5 Configuring Static Hosts on a VPRN SAP

The following displays a static host on a VPRN SAP:

```
*A:ALA-49>config>service>vprn>sub-if>grp-if# info
-----------------------------------------------
    description "VPRN service with static host"
    sap 7/1/5 create
        description "IES with static host"
        static-host ip 10.1.1.1 create
        static-host ip 2001::1/128 create
        static-host ip 2001:1::/64 create
    exit
-----------------------------------------------
*A:ALA-49>config>service>vprn#
```

2.11 Triple Play Services Command Reference

2.11.1 Command Hierarchies

Note:

- The command trees in this section are limited to those commands specific to Triple Play services. For the full command trees for a specific service type, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* and the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN*.
- Refer to the Triple Play DHCP Command Reference section for information about DHCP features, CLI syntax, and usage.

- IES/VPRN Triple Play Commands
- VPLS Triple Play Commands
- Show Commands
- Clear Commands
- Debug Commands
- Tools Commands

2.11.1.1 IES/VPRN Triple Play Commands

```
config
    — service
        — ies service-id [customer customer-id] [create] [vpn vpn-id] [name name]
        — no ies service-id
        — vprn service-id [name name] [customer customer-id] [create]
        — no vprn service-id
            — description description-string
            — no description
            — interface ip-int-name [create]
            — interface ip-int-name [create] tunnel
            — no interface ip-int-name
                — address {ip-address/mask | ip-address netmask} [broadcast all-ones
                    | host-ones] [track-srrp srrp-instance]
                — no address [ip-address/mask | ip-address netmask]
                — [no] allow-directed-broadcasts
                — [no] apr-learn-unsolicited
                — arp-limit limit [log-only] [threshold percent]
                — no arp-limit
```

- [no] arp-populate
- arp-timeout *seconds*
- no arp-timeout
- cflowd-parameters
  - sampling {unicast | multicast} type {acl | interface} [direction {ingress-only | egress-only | both}]
  - no sampling {unicast | multicast}
- description *description-string*
- no description
- host-connectivity-verify [source {vrrp | interface}] [interval *interval*] [action {remove | alarm}] [timeout *retry-timeout*] [retry-count *retry-count*]
- icmp
  - [no] mask-reply
  - param-problem [*number seconds*]
  - no param-problem
  - redirects [*number seconds*]
  - no redirects
  - ttl-expired [*number seconds*]
  - no ttl-expired
  - unreachables [*number seconds*]
  - no unreachables
- ip-mtu *octets*
- no ip-mtu
- [no] local-proxy-arp
- [no] loopback
- mac *ieee-address*
- no mac [*ieee-address*]
- proxy-arp-policy *policy-name* [*policy-name*]
- no proxy-arp-policy
- [no] remote-proxy-arp
- sap *sap-id* [create]
- no sap *sap-id*
  - accounting-policy *acct-policy-id*
  - no accounting-policy [*acct-policy-id*]
  - [no] collect-stats
  - description *description-string*
  - no description
  - egress
    - filter ip *ip-filter-id*
    - filter ipv6 *ipv6-filter-id*
    - no filter [ip *ip-filter-id*] [ipv6 *ipv6-filter-id*]
    - qos *policy-id*
    - qos *policy-id* port-redirect-group *queue-group-name* instance *instance-id*
    - no qos [policy-id]
    - [no] queue-override
      - queue *queue-id* [create]
      - no queue
        - adaptation-rule [pir *adaptation-rule*] [cir *adaptation-rule*]
        - avg-frame-overhead *percentage*
        - no avg-frame-overhead
        - cbs *size-in-kbytes*

```
                        — no cbs
                        — drop-tail low
                            — percent-reduction-from-mbs
                                percent
                            — no percent-reduction-from-mbs
                        — mbs size [bytes | kilobytes]
                        — no mbs
                        — rate pir-rate [cir cir-rate]
                        — no rate
                — [no] scheduler-override
                    — [no] scheduler scheduler-name
                        — rate pir-rate [cir cir-rate]
                        — no rate
                — scheduler-policy scheduler-policy-name
                — no scheduler-policy
            — ingress
                — filter ip ip-filter-id
                — filter ipv6 ipv6-filter-id
                — filter mac mac-filter-id
                — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id] [mac mac-filter-
                    id]
                — match-qinq-dot1p {top | bottom}
                — no match-qinq-dot1p
                — qos policy-id
                — qos policy-id [shared-queuing | multipoint-shared] fp-
                    redirect-group queue-group-name instance
                    instance-id
                — no qos
                — [no] queue-override
                    — queue queue-id [create]
                    — no queue queue-id
                        — adaptation-rule [pir adaptation-rule] [cir
                            adaptation-rule]
                        — avg-frame-overhead percentage
                        — no avg-frame-overhead
                        — cbs size-in-kbytes
                        — no cbs
                        — drop-tail
                            — percent-reduction-from-mbs
                                percent
                            — no percent-reduction-from-mbs
                        — mbs size [bytes | kilobytes]
                        — no mbs
                        — rate pir-rate [cir cir-rate]
                        — no rate
                — scheduler-policy scheduler-policy-name
                — no scheduler-policy
            — multi-service-site customer-site-name
            — no multi-service-site
            — static-host ip ip-address [mac ieee-address] [create]
            — static-host mac ieee-address [create]
            — no static-host [ip ip-address] mac ieee-address
            — no static-host all [force]
            — no static-host ip ip-address
```

— **ancp-string** *ancp-string*
— **no ancp-string**
— **app-profile** *app-profile-name*
— **no app-profile**
— **inter-dest-id** *intermediate-destination-id*
— **no inter-dest-id**
— **sla-profile** *sla-profile-name*
— **no sla-profile**
— **sub-profile** *sub-profile-name*
— **no sub-profile**
— **subscriber** *sub-ident*
— **no subscriber**
— [**no**] **subscriber-sap-id**
— [**no**] **shutdown**
— **spoke-sdp** *sdp-id:vc-id* [**vc-type** {**ether** | **ipipe**}] [**create**]
— **egress**
— **filter ip** *ip-filter-id*
— **filter ipv6** *ipv6-filter-id*
— **no filter** [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]
— **vc-label** *egress-vc-label*
— **no vc-label** [*egress-vc-label*]
— **ingress**
— **filter ip** *ip-filter-id*
— **filter ipv6** *ipv6-filter-id*
— **no filter** [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]
— **vc-label** *ingress-vc-label*
— **no vc-label** [*ingress-vc-label*]
— **secondary** {*ip-address/mask* | *ip-address netmask*} [**broadcast all-ones** | **host-ones**] [**igp-inhibit**]
— **no secondary** *ip-address*
— **static-arp** *ip-address ieee-address*
— **no static-arp** *ip-address* [*ieee-address*]
— [**no**] **shutdown**

2.11.1.2 VPLS Triple Play Commands

```
config
    — service
        — vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [b-vpls | i-vpls]
            [etree] [name name]
        — no vpls service-id
            — description description-string
            — no description
            — [no] disable-aging
            — [no] disable-learning
            — [no] discard-unknown
            — [no] fdb-table-high-wmark high-water-mark
            — [no] fdb-table-low-wmark low-water-mark
            — fdb-table-size table-size
            — no fdb-table-size [table-size]
            — local-age aging-timer
            — no local-age
            — mac-protect
                — [no] mac ieee-address
            — mac-subnet-length subnet-length
            — no mac-subnet-length
            — mcr-default-gtw
                — ip ip-address
                — no ip
                — mac ieee-address
                — no mac
            — remote-age seconds
            — no remote-age
            — service-mtu octets
            — no service-mtu
            — [no] shutdown
            — [no] split-horizon-group [group-name] [residential-group]
                — description description-string
                — no description
            — sap sap-id [split-horizon-group group-name] [create] [capture-sap] [eth-ring
                ring-index]
            — sap sap-id [split-horizon-group group-name] [create] [capture-sap] [eth-ring
                ring-index] leave-ac
            — sap sap-id [split-horizon-group group-name] [create] [capture-sap] [eth-ring
                ring-index] root-leaf-tag leaf-tag
            — no sap sap-id
                — accounting-policy acct-policy-id
                — no accounting-policy [acct-policy-id]
                — authentication-policy name
                — no authentication-policy
                — arp-host
                    — host-limit max-num-hosts
                    — no host-limit
                    — min-auth-interval min-auth-interval
                    — no min-auth-interval
                    — [no] shutdown
                — arp-reply-agent [sub-ident]
```

- **no arp-reply-agent**
- **bpdu-translation** {**auto** | **auto-rw** | **pvst** | **pvst-rw** | **stp**}
- **no bpdu-translation**
- **calling-station-id** *calling-station-id*
- **no calling-station-id**
- **description** *description-string*
- **no description**
- [**no**] **collect-stats**
- [**no**] **disable-aging**
- [**no**] **disable-learning**
- [**no**]**discard-unknown-source**
- **egress**
  - **filter ip** *ip-filter-id*
  - **filter ipv6** *ipv6-filter-id*
  - **no filter** [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]
  - **qos** *policy-id*
  - **qos** *policy-id* **port-redirect-group** *queue-group-name* **instance** **instance-id**
  - **no qos** [*policy-id*]
  - [**no**] **queue-override**
    - **queue** *queue-id* [**create**]
    - **no queue** *queue-id*
      - **adaptation-rule** [**pir** *adaptation-rule*] [**cir** *adaptation-rule*]
      - **avg-frame-overhead** *percentage*
      - **no avg-frame-overhead**
      - **cbs** *size-in-kbytes*
      - **no cbs**
      - **drop-tail**
        - **low**
          - **percent-reduction-from-mbs** *percent*
          - **no percent-reduction-from-mbs**
      - **mbs** *size* [**bytes** | **kilobytes**]
      - **no mbs**
      - **rate** *pir-rate* [**cir** *cir-rate*]
      - **no rate**
  - [**no**] **scheduler-override**
    - [**no**] **scheduler** *scheduler-name*
      - **rate** *pir-rate* [**cir** *cir-rate*]
      - **no rate**
  - **scheduler-policy** *scheduler-policy-name*
  - **no scheduler-policy**
- **ingress**
  - **filter ip** *ip-filter-id*
  - **filter ipv6** *ipv6-filter-id*
  - **filter mac** *mac-filter-id*
  - **no filter** [**ip** *ip-filter-id*] [**mac** *mac-filter-id*] [**ipv6** *ipv6-filter-id*]
  - **qos** *policy-id*
  - **qos** *policy-id* [**shared-queuing** | **multipoint-shared**] **fp-redirect-group** *queue-group-name* **instance** *instance-id*
  - **no qos** [*policy-id*]
  - [**no**] **queue-override**
    - [**no**] **queue** *queue-id*

— **adaptation-rule** [**pir** {**max** | **min** | **closest**}] [**cir** {**max** | **min** | **closest**}]
— **no adaptation-rule**
— **cbs** *size-in-kbytes*
— **no cbs**
— **drop-tail low**
— **percent-reduction-from-mbs** *percent*
— **no percent-reduction-from-mbs**
— **mbs** *size* [**bytes** | **kilobytes**]
— **no mbs**
— **rate** *pir-rate* [**cir** *cir-rate*]
— **no rate**
— [**no**] **scheduler-override**
— [**no**] **scheduler** *scheduler-name*
— **rate** *pir-rate* [**cir** *cir-rate*]
— **no rate**
— **scheduler-policy** *scheduler-policy-name*
— **no scheduler-policy**
— [**no**] **l2pt-termination**
— **max-nbr-mac-addr** *table-size*
— **no max-nbr-mac-addr**
— **multi-service-site** *customer-site-name*
— **no multi-service-site**
— **static-host ip** *ip-address* [**mac** *ieee-address*] [**create**]
— **static-host mac** *ieee-address* [**create**]
— **no static-host** [**ip** *ip-address*] **mac** *ieee-address*
— **no static-host all** [**force**]
— **no static-host ip** *ip-address*
— **ancp-string** *ancp-string*
— **no ancp-string**
— **app-profile** *app-profile-name*
— **no app-profile**
— **inter-dest-id** *intermediate-destination-id*
— **no inter-dest-id**
— **sla-profile** *sla-profile-name*
— **no sla-profile**
— **sub-profile** *sub-profile-name*
— **no sub-profile**
— **subscriber** *sub-ident*
— **no subscriber**
— [**no**] **subscriber-sap-id**
— [**no**] **static-mac** *ieee-mac-address*
— [**no**] **shutdown**
— **track-srrp** *srrp-id*
— **no track-srrp**
— **mesh-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**]
— **mesh-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**] **leaf-ac**
— **mesh-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**] **root-leaf-tag**
— **no mesh-sdp** *sdp-id*[:*vc-id*]
— **accounting-policy** *acct-policy-id*
— **no accounting-policy** [*acct-policy-id*]
— [**no**] **collect-stats**
— **egress**
— **filter ip** *ip-filter-id*

— **filter ipv6** *ipv6-filter-id*
— **filter mac** *mac-filter-id*
— **no filter** [**ip** *ip-filter-id*] [**mac** *mac-filter-id*] [**ipv6** *ipv6-filter-id*]
— **vc-label** *egress-vc-label*
— **no vc-label** [*egress-vc-label*]
— **ingress**
— **filter ip** *ip-filter-id*
— **filter ipv6** *ipv6-filter-id*
— **filter mac** *mac-filter-id*
— **no filter** [**ip** *ip-filter-id*] [**mac** *mac-filter-id*] [**ipv6** *ipv6-filter-id*]
— **vc-label** *ingress-vc-label*
— **no vc-label** [*ingress-vc-label*]
— [**no**] **shutdown**
— [**no**] **static-mac**
— **vlan-vc-tag** *vlan-id*
— **no vlan-vc-tag** [*vlan-id*]
— [**no**] **send-flush-on-failure**
— **spoke-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*] [**create**] **endpoint** *endpoint* **root-leaf-tag**
— **spoke-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*] **endpoint** *endpoint* **leaf-ac**
— **spoke-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*] [**create**] [**no-endpoint**]
— **spoke-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*] [**create**] **endpoint** *endpoint*
— **spoke-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*] [**create**] [**no-endpoint**] **leaf-ac**
— **spoke-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*] [**create**] [**no-endpoint**] **root-leaf-tag**
— **no spoke-sdp** *sdp-id*[:*vc-id*]
— **accounting-policy** *acct-policy-id*
— **no accounting-policy** [*acct-policy-id*]
— [**no**] **block-on-mesh-failure**
— [**no**] **collect-stats**
— **egress**
— **filter ip** *ip-filter-id*
— **filter ipv6** *ipv6-filter-id*
— **filter mac** *mac-filter-id*
— **no filter** [**ip** *ip-filter-id*] [**mac** *mac-filter-id*] [**ipv6** *ipv6-filter-id*]
— **vc-label** *egress-vc-label*
— **no vc-label** [*egress-vc-label*]
— **ingress**
— **filter ip** *ip-filter-id*
— **filter ipv6** *ipv6-filter-id*
— **filter mac** *mac-filter-id*
— **no filter** [**ip** *ip-filter-id*] [**mac** *mac-filter-id*] [**ipv6** *ipv6-filter-id*]
— **vc-label** *ingress-vc-label*
— **no vc-label** [*ingress-vc-label*]
— [**no**] **l2pt-termination**
— **max-nbr-mac-addr** *table-size*
— **no max-nbr-mac-addr**
— [**no**] **shutdown**
— [**no**] **static-mac** *ieee-mac-address*
— **vlan-vc-tag** *vlan-id*

— [**no**] **vlan-vc-tag** [*vlan-id*]

2.11.1.3 Show Commands

```
show
  — service
      — active-subscribers detail
      — active-subscribers mirror
      — active-subscribers [summary]
          — credit-control credit-control [subscriber sub-ident-string]
          — credit-control out-of-credit [action action] [summary]
          — filter [subscriber sub-ident-string] [origin origin]
          — hierarchy [subscriber sub-ident-string]
              — diameter
              — radius-acct
          — host-tracking [subscriber sub-ident-string]
          — host-tracking [subscriber sub-ident-string] detail
          — host-tracking [subscriber sub-ident-string] summary
          — host-tracking [subscriber sub-ident-string] statistics
              — groups [group group-ip-address]
              — groups group group-ip-address] detail
              — groups group group-ip-address] summary
          — igmp [subscriber sub-ident-string][detail]
          — subscriber sub-ident-string
          — subscriber sub-ident-string detail
          — subscriber sub-ident-string mirror
          — subscriber sub-ident-string sap sap-id sla-profile sla-profile-name [spi-
              sharing-type spi-sharing-type] [spi-sharing-id spi-sharing-id]
          — subscriber sub-ident-string sap sap-id sla-profile sla-profile-name [spi-
              sharing-type spi-sharing-type] [spi-sharing-id spi-sharing-id] detail
          — subscriber sub-ident-string sap sap-id sla-profile sla-profile-name [spi-
              sharing-type spi-sharing-type] [spi-sharing-id spi-sharing-id] mirror
      — id service-id
          — arp [ip-address] | [mac ieee-address] | [sap port-id:encap] | [interface ip-int-
              name]
          — base
          — authentication
              — statistics [policy name] [sap sap-id]
          — igmp-snooping
              — all
              — base
              — mrouters [detail]
              — mvr
              — port-db sap sap-id [detail]
              — port-db sap sap-id group grp-address
              — port-db sdp-id:vc-id [detail]
              — port-db sdp sdp-id:vc-id group grp-address
              — port-db {sap sap-id | sdp sdp-id:vc-id} [group grp-address] | detail]
              — proxy-db [detail]
              — proxy-db [group grp-address | detail]
              — querier
              — static [sap sap-id | sdp sdp-id:vc-id]
              — statistics [sap sap-id | sdp sdp-id:vc-id]
          — mfib [brief | group grp-address | statistics [group grp-address]]
          — mld-snooping
```

— **all**
— **base**
— **mrouters** [**detail**]
— **mvr**
— **port-db sap** *sap-id*
— **port-db sap** *sap-id* **detail**
— **port-db sap** *sap-id* **group** *grp-ipv6-address*
— **port-db sdp** *sdp-id:vc-id*
— **port-db sdp** *sdp-id:vc-id* **detail**
— **port-db sdp** *sdp-id:vc-id* **group** *grp-ipv6-address*
— **proxy-db** [**detail**]
— **proxy-db group** *grp-ipv6-address*
— **querier**
— **static** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]
— **statistics** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

— **retailers**
— **split-horizon-group** [*group-name*]
— **static-host** [**sap** *sap-id*] [**wholesaler** *service-id*] [**port** *port-id*] [**inter-dest-id** *intermediate-destination-id*] [**detail**] [**mac** *ieee-address*] [**ip-address***ip-prefix*[*/prefix-length*]]
— **static-host** [**sap** *sap-id*] [**wholesaler** *service-id*] [**port** *port-id*] **no-inter-dest-id** [**detail**]
— **static-host summary**
— **subscriber-hosts** [**sap** *sap-id*] [**ip** *ip-prefix/prefix-length*] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**app-profile** *app-profile-name*] [**wholesaler** *service-id*] [**address-origin** *address-origin*] [**detail**] [**statistics**] [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]
— **wholesalers**

— **subscriber-using** [**service-id** *service-id*] [**sap-id** *sap-id*] [**interface** *ip-int-name*] [**ip** *ip-address*[*/mask*]] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**app-profile** *app-profile-name*] [**port** *port-id*] [**no-inter-dest-id** | **inter-dest-id** *intermediate-destination-id*]

show
— **redundancy**
— **multi-chassis all**
— **multi-chassis mc-lag**
— **multi-chassis sync**
— **mc-ipsec** *addr* [**tunnel-group** *grp-id*]
— **mc-ring peer** *ip-address* **statistics**
— **mc-ring peer** *ip-address* [**ring** *sync-tag* [**detail** | **statistics**]]
— **mc-ring peer** *ip-address* **ring** *sync-tag* **ring-node** [*node-name* [**detail** | **statistics**]]
— **mc-ring global-statistics**

show
— **router**
— **igmp**
— **group** [*grp-ip-address*]
— **group summary**
— **group-interface** [**fwd-service** *service-id*] [*ip-int-name*] [**detail**]
— **hosts** [**group** *grp-address*] [**detail**] [**fwd-service** *service-id*] [**grp-interface** *ip-int-name*]

```
            — hosts [host ip-address] [group grp-address] [detail]
            — hosts summary
            — interface [ip-int-name | ip-address] [group] [grp-ip-address] [detail]
            — mcast-reporting-statistics [host]
            — mld
                — interface port:tags
            — ssm-translate [interface-name]
            — static [ip-int-name | ip-addr]
            — statistics [ip-int-name | ip-address]
            — statistics host [ip-address]
            — status
            — tunnel-interface
        — pim
            — mc-ecmp-balance [detail]
```

2.11.1.4 Clear Commands

```
clear
    — service
        — id service-id
            — arp-host [wholesaler service-id] [sap sap-id | interface interface-name | ip-
              address ip-address[/mask] | mac ieee-address|{[port port-id] [no-inter-
              dest-id | inter-dest-id inter-dest-id]}] [detail]
            — arp-host statistics [sap sap-id | interface interface-name]
            — arp-host summary [ interface interface-name | saps ]
            — authentication
                — statistics
            — igmp-snooping
                — port-db sap sap-id [group grp-address [source ip-address]]
                — port-db sdp sdp-id:vc-id [group grp-address [source ip-address]]
                — querier
                — querier [all | sap sap-id | sdp sdp-id:vc-id
                — statistics [all | sap sap-id | sdp sdp-id:vc-id]
            — mfib
                — statistics [all | group grp-address]
            — mld-snooping
                — port-db sap sap-id [group grp-ipv6-address]
                — port-db sap sap-id group grp-ipv6-address source src-ipv6-address
                — port-db sdp sdp-id:vc-id [group grp-ipv6-address]
                — port-db sdp sdp-id:vc-id group grp-ipv6-address source src-ipv6-
                  address
                — querier
                — statistics all
                — statistics sap sap-id
                — statistics sdp sdp-id:vc-id
            — msap {msap-id | idle-only}
            — msap-policy msap-policy-name
```

2.11.1.5 Debug Commands

```
debug
    — service
        — [no] id service-id
            — [no] arp-host
            — [no] mld-snooping
                — detail-level {low | medium | high}
                — no detail-level
                — [no] mac ieee-address
                — mode {dropped-only | ingr-and-dropped | egr-ingr-and-dropped}
                — no mode
                — [no] sap sap-id
                — [no] sdp sdp-id:vc-id
            — detail-level {low | medium | high}
            — no detail-level
            — [no] mac ieee-address
            — mode {dropped-only | ingr-and-dropped | egr-ingr-and-dropped}
            — no mode
            — [no] sap sap-id
            — [no] sdp sdp-id:vc-id
```

2.11.1.6 Tools Commands

```
    — tools
        — dump
            — service
                — msap [service service-id]
```

2.11.2 Command Descriptions

2.11.2.1 Global Commands

description

Syntax **description** *description-string*
no description

Context config>service>vpls
config>service>vpls>split-horizon-group

config>service>vpls>sap
config>service>ies>if>sap
config>service>vprn

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Parameters *description-string* — Specifies the description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

shutdown

Syntax [**no**] **shutdown**

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>sap>arp-host
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp
config>service>ies>if>sap
config>service>vprn

Description This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

Services are created in the administratively down (**shutdown**) state. When a **no shutdown** command is entered, the service becomes administratively up and then tries to enter the operationally up state. Default administrative states for services and service entities is described below in Special Cases.

The **no** form of this command places the entity into an administratively enabled state.

Special Cases **Service Admin State —** Bindings to an SDP within the service are put into the out-of-service state when the service is shutdown. While the service is shutdown, all customer packets are dropped and counted as discards for billing and debugging purposes.

Service Operational State — A service is regarded as operational providing that two SAPs or if one SDP are operational.

SDP (global) — When an SDP is shutdown at the global service level, all bindings to that SDP are put into the out-of-service state and the SDP itself is put into the administratively and operationally down states. Packets that would normally be transmitted using this SDP binding are discarded and counted as dropped packets.

SDP (service level) — Shutting down an SDP within a service only affects traffic on that service from entering or being received from the SDP. The SDP itself may still be operationally up for other services.

SDP Keepalives — Enables SDP connectivity monitoring keepalive messages for the SDP ID. Default state is disabled (shutdown) in which case the operational state of the SDP-ID is not affected by the keepalive message state.

VPLS SAPs and SDPs — SAPs are created in a VPLS and SDPs are bound to a VPLS in the administratively up default state. The created SAP attempts to enter the operationally up state. An SDP attempts to go into the in-service state once bound to the VPLS.

2.11.2.2 IES/VPRN Service Commands

ies

Syntax **ies** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**name** *name*]
no ies *service-id*

Context config>service

Description This command creates or edits an IES service instance.

The **ies** command creates or maintains an Internet Ethernet Service (IES). If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

IES services allow the creation of customer facing IP interfaces in the same routing instance used for service network core routing connectivity. IES services require that the IP addressing scheme used by the subscriber must be unique between it and other addressing schemes used by the provider and potentially the entire Internet.

While IES is part of the routing domain, the usable IP address space may be limited. This allows a portion of the service provider address space to be set aside for service IP provisioning, becoming administered by a separate but subordinate address authority. This feature is defined using the **config router service-prefix** command.

IP interfaces defined within the context of an IES service ID must have a SAP created as the access point to the subscriber network. This allows a combination of bridging and IP routing for redundancy purposes.

When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the **customer** command in the service context. Once a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and recreated with a new customer association.

Once a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified results in an error.

Multiple IES services are created to separate customer owned IP interfaces. More than one IES service may be created for a single customer ID. More than one IP interface may be created within a single IES service ID. All IP interfaces created within an IES service ID belongs to the same customer.

By default, no IES service instances exist until they are explicitly created.

The **no** form of this command deletes the IES service instance with the specified *service-id*. The service cannot be deleted until all the IP interfaces defined within the service ID have been shutdown and deleted.

Parameters *service-id* — Specifies the unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every router on which this service is defined.

Values *service-id*: 1 to 214748364

svc-name: A string up to 64 characters

customer-id — Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn-id — Specifies the VPN ID number used to identify virtual private networks (VPNs) by a VPN identification number.

Values 1 to 2147483647

Default **null** (0)

create — Keyword used to create the service ID. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

name *name* — This parameter configures an optional service name, up to 64 characters, which adds a name identifier to a given service to then use that service name in configuration references as well as display and use service names in show commands throughout the system. This helps the service provider/administrator to identify and manage services within the SR OS platforms.

All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used to identify and reference a given service once it is initially created.

If a name is not specified at creation time, then SR OS assigns a string version of the service-id as the name.

Service names may not begin with an integer (0 to 9).

Values *name*: up to 64 characters

vprn

Syntax **vprn** *service-id* [**name** *name*] [**customer** *customer-id*] [**create**]
no vprn *service-id*

Context config>service

Description This command creates or edits a Virtual Private Routed Network (VPRN) service instance.

If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

VPRN services allow the creation of customer-facing IP interfaces in the same routing instance used for service network core routing connectivity. VPRN services require that the IP addressing scheme used by the subscriber must be unique between it and other addressing schemes used by the provider and potentially the entire Internet.

IP interfaces defined within the context of an VPRN service ID must have a SAP created as the access point to the subscriber network.

When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the customer command in the service context. When a service is created with a customer association, it is not possible to edit the customer association. The service must be deleted and re-created with a new customer association.

When a service is created, the use of the **customer** *customer-id* is optional to navigate into the service configuration context. If attempting to edit a service with the incorrect *customer-id* results in an error.

Multiple VPRN services are created to separate customer-owned IP interfaces. More than one VPRN service can be created for a single customer ID. More than one IP interface can be created within a single VPRN service ID. All IP interfaces created within an VPRN service ID belongs to the same customer.

The **no** form of this command deletes the VPRN service instance with the specified *service-id*. The service cannot be deleted until all the IP interfaces and all routing protocol configurations defined within the service ID have been shutdown and deleted.

Parameters *service-id* — Specifies the unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7750 SR on which this service is defined.

Values

service-id:	1 to 2147483648
svc-name:	64 characters maximum

customer-id — Specifies an existing customer identification number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

name *name* — This parameter configures an optional vprn name, up to 64 characters, which adds a name identifier to a given vprn to then use that vprn name in configuration references as well as display and use vprn names in show commands throughout the system. This helps the service provider/administrator to identify and manage vprn within the SR OS platforms.

All services are required to assign a service ID to initially create a service. However, either the service ID or the service name can be used to identify and reference a given service once it is initially created.

If a name is not specified at creation time, then SR OS assigns a string version of the service-id as the name.

Service names may not begin with an integer (0 to 9).

Values *name*:64 characters maximum

create — Keyword used to create a service ID. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

vpls

Syntax **vpls** *service-id* [**customer** *customer-id*] [**vpn** *vpn-id*] [**m-vpls**] [**b-vpls** | **i-vpls**] [**etree**] [**name** *name*] [**create**]
no vpls *service-id*

Context config>service

Description This command creates or edits a Virtual Private LAN Services (VPLS) instance. The **vpls** command is used to create or maintain a VPLS service. If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

A VPLS service connects multiple customer sites together acting like a zero-hop, layer 2 switched domain. A VPLS is always a logical full mesh.

When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the **customer** command in the service context. Once a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and recreated with a new customer association.

To create a management VPLS, the **m-vpls** keyword must be specified. See section **Hierarchical VPLS Redundancy** for an introduction to the concept of management VPLS.

Once a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified results in an error.

More than one VPLS service may be created for a single customer ID.

By default, no VPLS instances exist until they are explicitly created.

The **no** form of this command deletes the VPLS service instance with the specified *service-id*. The service cannot be deleted until all SAPs and SDPs defined within the service ID have been shutdown and deleted, and the service has been shutdown.

Parameters *service-id* — Specifies the unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every router on which this service is defined.

Values *service-id*: 1 to 214748364

svc-name: A string up to 64 characters.

customer-id — Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn-id — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number.

Values 1 to 2147483647

Default **null** (0)

m-vpls — Specifies a managed VPLS.

name *name* — Configures an optional service name identifier, up to 64 characters, to a given service. This service name can then be used in configuration references, display, and show commands throughout the system. A defined service name can help the service provider or administrator to identify and manage services within the SR OS platforms.

To create a service, you must assign a service ID; however, after it is created, either the service ID or the service name can be used to identify and reference a service.

If a name is not specified at creation time, then SR OS assigns a string version of the *service-id* as the name.

Service names may not begin with an integer (0 to 9).

Values *name*: 64 characters maximum

create — Keyword used to create the service ID. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

interface

Syntax **interface** *ip-int-name* [**create**]
interface *ip-int-name* [**create**] **tunnel**
no interface *ip-int-name*

Context config>service>ies
config>service>vprn

Description This command creates a logical IP routing interface for a Virtual Private Routed Network (VPRN). Once created, attributes like an IP address and service access point (SAP) can be associated with the IP interface.

The **interface** command, under the context of services, is used to create and maintain IP routing interfaces within VPRN service IDs. The **interface** command can be executed in the context of an VPRN service ID. The IP interface created is associated with the service core network routing instance and default routing table. The typical use for IP interfaces created in this manner is for subscriber Internet access.

Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for **config router interface** and **config service vprn interface**. Interface names must not be in the dotted decimal notation of an IP address. For example, the name “1.1.1.1” is not allowed, but “int-1.1.1.1” is allowed. Show commands for router interfaces use either interface names or the IP addresses. Use unique IP address values and IP address names to maintain clarity. It could be unclear to the user if the same IP address and IP address name values are used. Although not recommended, duplicate interface names can exist in different router instances.

The available IP address space for local subnets and routes is controlled with the **config router service-prefix** command. The **service-prefix** command administers the allowed subnets that can be defined on service IP interfaces. It also controls the prefixes that may be learned or statically defined with the service IP interface as the egress interface. This allows segmenting the IP address space into **config router** and **config service** domains.

When a new name is entered, a new logical router interface is created. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, there are no default IP interface names defined within the system. All VPRN IP interfaces must be explicitly defined. Interfaces are created in an enabled state.

The **no** form of this command removes IP the interface and all the associated configuration. The interface must be administratively shutdown before issuing the **no interface** command.

For VPRN services, the IP interface must be shutdown before the SAP on that interface may be removed. VPRN services do not have the **shutdown** command in the SAP CLI context. VPRN service SAPs rely on the interface status to enable and disable them.

Parameters *ip-int-name* — Specifies the name of the IP interface. Interface names must be unique within the group of defined IP interfaces for **config router interface** and **config service vprn interface** commands. An interface name cannot be in the form of an IP address. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

If *ip-int-name* already exists within the service ID, the context is changed to maintain that IP interface. If *ip-int-name* already exists within another service ID or is an IP interface defined within the **config router** commands, an error occurs and context is changed to that IP interface. If *ip-int-name* does not exist, the interface is created and context is changed to that interface for further command processing.

tunnel — Specifies that this is an IPsec interface used for IPsec tunneling.

create — Keyword used to create the interface. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

address

Syntax **address** {*ip-address/mask* | *ip-address netmask*} [**broadcast all-ones** | **host-ones**] [**track-srrp** *srrp-instance*]
no address [*ip-address/mask* | *ip-address netmask*]

Context config>service>ies>if
config>service>vprn>if

Description This command assigns an IP address, IP subnet, and broadcast address format to an IES IP router interface. Only one IP address can be associated with an IP interface.

An IP address must be assigned to each IES IP interface. An IP address and a mask are used together to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the router.

The local subnet that the **address** command defines must be part of the services' address space within the routing context using the **config router service-prefix** command. The default is to disallow the complete address space to services. Once a portion of the address space is allocated as a service prefix, that portion can be made unavailable for IP interfaces defined within the **config router interface** CLI context for network core connectivity with the **exclude** option in the **config router service-prefix** command.

The IP address for the interface can be entered in either CIDR (Classless Inter-Domain Routing) or traditional dotted decimal notation. The show commands display CIDR notation and is stored in configuration files.

By default, no IP address or subnet association exists on an IP interface until it is explicitly created.

Table 7 **Address Admin and Operational States**

Address	Admin State	Oper State
No address	up	down
No address	down	down
1.1.1.1	up	up
1.1.1.1	down	down

The operational state is a read-only variable and the only controlling variables are the address and admin states. The address and admin states are independent and can be set independently. If an interface is in an administratively up state and an address is assigned, it becomes operationally up and the protocol interfaces and the MPLS LSPs associated with that IP interface are reinitialized.

The **no** form of this command removes the IP address assignment from the IP interface. When the **no address** command is entered, the interface becomes operationally down.

Parameters *ip-address* — Specifies the IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that is used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

/ — Specifies a parameter delimiter and separates the *ip-address* portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the "*/*" and the *mask-length* parameter. If a forward slash is not immediately following the *ip-address*, a dotted decimal mask must follow the prefix.

mask-length — Specifies the subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-address* from the *mask-length* parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address.

Note: A mask length of 32 is reserved for loopback addresses (includes system addresses).

Default 0 to 31

mask — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that is used in a logical 'AND' function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252.

Note: A mask of 255.255.255.255 is reserved for system IP addresses.

broadcast — Overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is **host-ones** which indicates a subnet broadcast address. Use this parameter to change the broadcast address to **all-ones** or revert back to a broadcast address of **host-ones**.

The broadcast format on an IP interface can be specified when the IP address is assigned or changed.

This parameter does not affect the type of broadcasts that can be received by the IP interface. A host sending either the local broadcast (**all-ones**) or the valid subnet broadcast address (**host-ones**) is received by the IP interface.

Default host-ones

all-ones — Specifies the broadcast address used by the IP interface for this IP address is 255.255.255.255, also known as the local broadcast.

host-ones — Specifies that the broadcast address used by the IP interface for this IP address is the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the *ip-address* and the *mask-length* or *mask* with all the host bits set to binary one. This is the default broadcast address used by an IP interface.

The **broadcast** parameter within the **address** command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the **broadcast** type to **host-ones** after being changed to **all-ones**, the **address** command must be executed with the **broadcast** parameter defined.

allow-directed-broadcasts

Syntax [**no**] **allow-directed-broadcasts**

Context config>service>ies>if
config>service>vprn>if

Description This command enables the forwarding of directed broadcasts out of the IP interface.

A directed broadcast is a packet received on a local router interface destined for the subnet broadcast address on another IP interface. The **allow-directed-broadcasts** command on an IP interface enables or disables the transmission of packets destined to the subnet broadcast address of the egress IP interface.

When enabled, a frame destined to the local subnet on this IP interface is sent as a subnet broadcast out this interface. Care should be exercised when allowing directed broadcasts as it is a well-known mechanism used for denial-of-service attacks.

When disabled, directed broadcast packets discarded at this egress IP interface are counted in the normal discard counters for the egress SAP.

By default, directed broadcasts are not allowed and are discarded at this egress IP interface.

The **no** form of this command disables the forwarding of directed broadcasts out of the IP interface. All broadcasts are dropped.

apr-learn-unsolicited

Syntax [**no**] **arp-learn-unsolicited**

Context config>service>ies>interface
config>service>vprn>interface

Description This command enables ARP learn unsolicited.

arp-limit

Syntax **arp-limit** *limit* [**log-only**] [**threshold** *percent*]
no arp-limit

Context config>service>ies>interface

Description This command configures the maximum amount of dynamic IPv4 ARP entries that can be learned on an IP interface.

When the number of dynamic ARP entries reaches the configured percentage of this limit, an SNMP trap is sent. When the limit is exceeded, no new entries are learned until an entry expires and traffic to these destinations is dropped. Entries that have already been learned are refreshed.

The **no** form of this command removes the **arp-limit**.

Default arp-limit threshold 90

Parameters **log-only** — Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, entries above the limit are learned.

percent — The threshold value (as a percentage) that triggers a warning message to be sent.

Values 0 to 100

limit — The number of entries that can be learned on an IP interface expressed as a decimal integer. If the limit is set to 0, dynamic ARP learning is disabled and no dynamic ARP entries are learned.

Values 0 to 524288

arp-populate

Syntax [**no**] **arp-populate**

Context config>service>ies>if
config>service>vprn>if

Description This command, when enabled, disables dynamic learning of ARP entries. Instead, the ARP table is populated with dynamic entries from the DHCP Lease State Table (enabled with **lease-populate**), and optionally with static entries entered with the **host** command.

Enabling the **arp-populate** command removes any dynamic ARP entries learned on this interface from the ARP cache.

The **arp-populate** command fails if an existing static ARP entry exists for this interface.

The **arp-populate** command fails if an existing static subscriber host on the SAP does not have both MAC and IP addresses specified.

Once **arp-populate** is enabled, creating a static subscriber host on the SAP without both an IP address and MAC address fails.

When **arp-populate** is enabled, the system does not send out ARP requests for hosts that are not in the ARP cache. Only statically configured and DHCP learned hosts are reachable through an IP interface with **arp-populate** enabled. The **arp-populate** command can only be enabled on IES and VPRN interfaces supporting Ethernet encapsulation.

The **no** form of this command disables ARP cache population functions for static and dynamic hosts on the interface. All static and dynamic host information for this interface is removed from the system's ARP cache.

arp-timeout

Syntax **arp-timeout** *seconds*
no arp-timeout

Context config>service>ies>if
config>service>vprn>if

Description This command configures the minimum time, in seconds, an ARP entry learned on the IP interface is stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

When the **arp-populate** and **lease-populate** commands are enabled on an IES interface, the ARP table entries are no longer dynamically learned, but instead by snooping DHCP ACK message from a DHCP server. In this case the configured **arp-timeout** value has no effect.

The default value for **arp-timeout** is 14400 seconds (4 hours).

The **no** form of this command restores **arp-timeout** to the default value.

Default arp-timeout 14400

Parameters *seconds* — Specifies the minimum number of seconds a learned ARP entry is stored in the ARP table. A value of zero specifies that the timer is inoperative and learned ARP entries are not be aged.

Values 0 to 65535

cflowd-parameters

Syntax **cflowd-parameters**

Context config>service>ies>if

Description This command enables Cflowd to collect traffic flow samples through a router for analysis.

Cflowd is used for network planning and traffic engineering, capacity planning, security, application and user profiling, performance monitoring, usage-based billing, and SLA measurement. When Cflowd is enabled at the interface level, all packets forwarded by the interface are subjected to analysis according to the **cflowd** configuration.

sampling

Syntax **sampling** {**unicast** | **multicast**} **type** {**acl** | **interface**} [**direction** {**ingress-only** | **egress-only | both}]**
no sampling {**unicast** | **multicast**}

Context config>service>ies>if>cflowd-parameters

Description This command configures the Cflowd sampling parameters.

Parameters **unicast** — Specifies that the **sampling** command controls the sampling of unicast traffic on the associated interface/SAP.

multicast — Specifies that the **sampling** command controls the sampling of multicast traffic on the associated interface/SAP.

type — Specifies the cflowd sampling type on the specified virtual router interfaces.

Values acl — Specifies that the sampled traffic is controlled via an IP traffic filter entry with the action “filter-sample” configured.

interface — Specifies that all traffic entering or exiting the interface is subject to sampling.

direction — Specifies the direction in which to collect traffic flow samples.

Values ingress-only — Enables ingress sampling only on the associated interface.

egress-only — Enables egress sampling only on the associated interface.

both — Enables both ingress and egress cflowd sampling.

host-connectivity-verify

Syntax **host-connectivity-verify** [**source** {**vrrp** | **interface**}] [**interval** *interval*] [**action** {**remove** | **alarm**}] [**timeout** *retry-timeout*] [**retry-count** *retry-count*]

Context config>service>ies>if
config>service>vprn>if

Description This command enables subscriber host connectivity verification for all hosts on this interface. This tool periodically scans all known hosts (from dhcp-state) and perform UC ARP requests. The subscriber host connectivity verification maintains state (connected vs. not-connected) for all hosts.

The **no** form of this command reverts to the default.

Parameters **source** {**vrrp** | **interface**} — Specifies the source to be used for generation of subscriber host connectivity verification packets. The **vrrp** keyword specifies that the VRRP state should be used to select proper IP and MAC (active uses VRID, back-up uses interface addresses). The **interface** keyword forces the use of the interface mac and ip addresses.

Note: There are up to 256 possible subnets on a given interface, therefore, the subscriber host connectivity verification tool always uses an address of the subnet to which the given host is pertaining. For group-interfaces, one of the parent subscriber interface subnets (depending on host's address) is used.

action {**remove** | **alarm**} — Defines the action taken on a subscriber host connectivity verification failure for a given host. The **remove** keyword raises an alarm and removes dhcp-state and releases all allocated resources (queues, table entries, and so on). The **alarm** keyword raises an alarm indicating that the host is disconnected.

interval — Specifies the interval, expressed in minutes, which specifies when all known sources should be verified. The actual rate is then dependent on number of known hosts and interval.

Values 1 to 6000

Note: A zero value can be used by the SNMP agent to disable host-connectivity-verification.

retry-timeout — Specifies the timeout, in seconds, before a connectivity check retransmission.

Values 10 to 60

retry-count — Specifies the number of connectivity check retransmissions.

Values 2 to 29

icmp

Syntax **icmp**

Context config>service>ies>if
config>service>vprn>if

Description This command enables the context to configure Internet Control Message Protocol (ICMP) parameters on a service.

mask-reply

Syntax [**no**] **mask-reply**

Context config>service>ies>if>icmp
config>service>vprn>if

Description This command enables responses to Internet Control Message Protocol (ICMP) mask requests on the router interface.

If a local node sends an ICMP mask request to the router interface, the **mask-reply** command configures the router interface to reply to the request.

By default, the router instance replies to mask requests.

The **no** form of this command disables replies to ICMP mask requests on the router interface.

Default mask-reply

param-problem

Syntax **param-problem** [*number seconds*]
no param-problem

Context config>service>ies>if>ipv6>icmp6

Description This command specifies whether parameter-problem ICMPv6 messages should be sent. When enabled, parameter-problem ICMPv6 messages are generated by this interface.

The **no** form of this command disables the sending of parameter-problem ICMPv6 messages.

Default param-problem 100 10

Parameters *number* — Specifies the number of parameter-problem ICMPv6 messages to send in the time frame specified by the *seconds* parameter.

Values 10 to 1000

Default 100

seconds — Specifies the time frame, in seconds, that is used to limit the number of parameter-problem ICMPv6 messages issued.

Values 1 to 60

Default 10

redirects

Syntax **redirects** [*number seconds*]
no redirects

Context config>service>ies>if>icmp
config>service>vprn>if

Description This command configures the rate for Internet Control Message Protocol (ICMP) redirect messages issued on the router interface.

When routes are not optimal on this router and another router on the same subnetwork has a better route, the router can issue an ICMP redirect to alert the sending node that a better route is available.

The **redirects** command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects is issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of redirect messages that can be issued on the interface for a given time interval.

The **no** form of this command disables the generation of icmp redirects on the router interface.

Default redirects 100 10

Parameters *number* — Specifies the maximum number of ICMP redirect messages to send. This parameter must be specified with the *second* parameter.

Values 10 to 1000

seconds — Specifies the time frame in seconds used to limit the *number* of ICMP redirect messages that can be issued.

Values 1 to 60

ttl-expired

Syntax **ttl-expired [***number seconds***]**
no ttl-expired

Context config>service>ies>if>icmp

Description This command configures the rate Internet Control Message Protocol. (ICMP) TTL expired messages are issued by the IP interface.

By default, the generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the limiting the rate of TTL expired messages on the router interface.

Default ttl-expired 100 10

Parameters *number* — Specifies the maximum number of ICMP TTL expired messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 to 2000

seconds — Specifies the time frame, in seconds, used to limit the number of ICMP TTL expired messages that can be issued, expressed as a decimal integer.

Values 1 to 60

unreachables

Syntax **unreachables** [*number seconds*]
no unreachables

Context config>service>ies>if>icmp

Description This command configures the rate for ICMP host and network destination unreachable messages issued on the router interface.

The **unreachables** command enables the generation of ICMP destination unreachables on the router interface. The rate at which ICMP unreachables is issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of destination unreachable messages which can be issued on the interface for a given time interval.
By default, generation of ICMP destination unreachables messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the generation of icmp destination unreachables on the router interface.

Default unreachables 100 10

Parameters *number* — Specifies the maximum number of ICMP unreachable messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 to 2000

seconds — Specifies the time frame, in seconds, used to limit the *number* of ICMP unreachable messages that can be issued.

Values 1 to 60

ip-mtu

Syntax **ip-mtu** *octets*
no ip-mtu

Context config>service>ies>if
config>service>vprn>if

Description This command specifies the maximum size of IP packets on this group interface. Packets larger than this are fragmented.

The **ip-mtu** applies to all IPoE host types (dhcp, arp, static). For PPP/L2TP sessions, the **ip-mtu** is not taken into account for the MTU negotiation. The **ppp-mtu** in the **ppp-policy** should be used instead.

The **no** form of this command reverts to the default.

Parameters *octets* — Specifies the largest frame size (in octets) that this interface can handle.

Values 512 to 9000

local-proxy-arp

Syntax [**no**] **local-proxy-arp**

Context config>service>ies>if
config>service>vprn>if

Description This command enables local proxy ARP. When local proxy ARP is enabled on an IP interface, the system responds to all ARP requests for IP addresses belonging to the subnet with its own MAC address, and thus becomes the forwarding point for all traffic between hosts in that subnet.

When **local-proxy-arp** is enabled, ICMP redirects on the ports associated with the service are automatically blocked.

The **no** form of this command reverts to the default.

loopback

Syntax [**no**] **loopback**

Context config>service>ies>if
config>service>vprn>if

Description This command specifies that the associated interface is a loopback interface that has no associated physical interface. As a result, the associated IES/VPRN interface cannot be bound to a SAP.

Note: Configure an IES interface as a loopback interface by issuing the **loopback** command instead of the **sap** *sap-id* command. The loopback flag cannot be set on an interface where a SAP is already defined and a SAP cannot be defined on a loopback interface.

The **no** form of this command reverts to the default.

mac

Syntax **mac** *ieee-address*
no mac [*ieee-address*]

Context config>service>ies>if

Description This command assigns a specific MAC address to an IES IP interface.

The **no** form of this command returns the MAC address of the IP interface to the default value.

Default The physical MAC address associated with the Ethernet interface that the SAP is configured on (the default MAC address assigned to the interface, assigned by the system).

Parameters *ieee-address* — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

proxy-arp-policy

Syntax **proxy-arp-policy** *policy-name* [*policy-name*]
no proxy-arp-policy

Context config>service>ies>if

Description This command configures a proxy ARP policy for the interface.

The **no** form of this command disables the proxy ARP capability.

Parameters *policy-name* — Specifies up to five export route policy names. Allowed values are any string up to 32 characters composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

The specified name(s) must already be defined.

remote-proxy-arp

Syntax [**no**] **remote-proxy-arp**

Context config>service>ies>if
config>service>vprn>if

Description This command enables remote proxy ARP on the interface.

Remote proxy ARP is similar to proxy ARP. It allows the router to answer an ARP request on an interface for a subnet that is not provisioned on that interface. This allows the router to forward to the other subnet on behalf of the requester. To distinguish remote proxy ARP from local proxy ARP, local proxy ARP performs a similar function but only when the requested IP is on the receiving interface.

The **no** form of this command reverts to the default.

sap

Syntax **sap** *sap-id* [**create**]
no sap *sap-id*

Context config>service>ies>if
config>service>vprn>if

Description This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the router. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, then the SAP does not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **config interface** *port-type port-id* **mode access** command.

If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service is discarded. The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

Note: Configure an IES interface as a loopback interface by issuing the **loopback** command instead of the **sap** *sap-id* command. The loopback flag cannot be set on an interface where a SAP is already defined and a SAP cannot be defined on a loopback interface.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP are also deleted. For IES, the IP interface must be shutdown before the SAP on that interface may be removed.

Special Cases

IES — IES SAP can be defined with Ethernet or SONET/SDH ports on the 7450 ESS. A SAP is defined within the context of an IP routed interface. Each IP interface is limited to a single SAP definition. Attempts to create a second SAP on an IP interface fails and generates an error; the original SAP is not affected.

Parameters

sap-id — Specifies the SAP ID.

create — Keyword used to create the SAP. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

accounting-policy

Syntax

accounting-policy *acct-policy-id*
no accounting-policy

Context

config>service>ies>if>sap
config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vprn>if>sap

Description

This command creates the accounting policy context that can be applied to a SAP or SDP.

An accounting policy must be defined before it can be associated with a SAP or SDP.

If the *policy-id* does not exist, an error message is generated.

A maximum of one accounting policy can be associated with a SAP or SDP at one time. Accounting policies are configured in the **config>log** context.

The **no** form of this command removes the accounting policy association from the SAP or SDP, and the accounting policy reverts to the default.

Default Default accounting policy.

Parameters *acct-policy-id* — Specifies the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1 to 99

calling-station-id

Syntax [**no**] **calling-station-id** *calling-station-id*

Context config>service>ies>if>sap
config>service>vpls>sap
config>service>vprn>if
config>subscr-mgmt>auth-plcy>include-radius-attribute

Description This command enables the inclusion of the **calling-station-id** attribute in RADIUS authentication requests and RADIUS accounting messages.

The **no** form of this command reverts to the default.

Parameters *calling-station-id* — Specifies the calling station ID, up to 64 characters.

collect-stats

Syntax [**no**] **collect-stats**

Context config>service>ies>if>sap
config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp

Description This command enables accounting and statistical data collection for either the SAP or SDP, network port, or IP interface. When applying accounting policies the data, by default, is collected in the appropriate records and written to the designated billing file.

When the **no collect-stats** command is issued the statistics are still accumulated by the IOM cards. However, the CPU does not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default collect-stats

egress

Syntax **egress**

Context config>service>vpls>sap
config>service>vpls>mesh-sdp
config>service>ies>if>sap

Description This command enables the context to configure egress Quality of Service (QoS) policies and filter policies.

If no QoS policy is defined, the system default QoS policy is used for egress processing. If no egress filter is defined, no filtering is performed.

filter

Syntax **filter ip** *ip-filter-id*
filter ipv6 *ipv6-filter-id*
filter mac *mac-filter-id*
no filter [**ip** *ip-filter-id*] [**ipv6** *ipv6-filter-id*]

Context config>service>ies>if>sap>egress
config>service>ies>if>sap>ingress

Description This command associates an IP filter policy with an ingress or egress Service Access Point (SAP). Filter policies control the forwarding and dropping of packets based on the matching criteria. MAC filters are only allowed on Epipe and Virtual Private LAN Service (VPLS) SAPs.

The **filter** command is used to associate a filter policy with a specified *filter ID* with an ingress or egress SAP. The *filter policy* must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation fails and an error message returned.

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to the match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

Special Cases **IES —** Only IP filters are supported on an IES IP interface, and the filters only apply to routed traffic.

Parameters *ip-filter-id* — Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 to 65535

ipv6-filter-id — Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.

Values 1 to 65535

qos

Syntax **qos** *policy-id*
qos *policy-id* **port-redirect-group** *queue-group-name* **instance** *instance-id*
no qos [**policy-id**]

Context config>service>vpls>sap>egress
config>service>ies>if>sap>egress
config>service>vprn>if>sap>egress

Description This command associates a Quality of Service (QoS) policy with an egress Service Access Point (SAP) or IP interface.

QoS egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP or IP interface. If the policy ID does not exist, an error is returned.

The **qos** command associates both ingress and egress QoS policies. The **qos** command only allows ingress policies to be associated on SAP or IP interface ingress and egress policies on SAP or IP interface egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type returns an error.

By default, no specific QoS policy is associated with the SAP or IP interface for ingress or egress, so the default QoS policy is used.

The **no** form of this command removes the QoS policy association from the SAP or IP interface, and the QoS policy reverts to the default.

Parameters *policy-id* — Specifies ingress or egress policy ID to associate with SAP or IP interface on ingress or egress. The policy ID must already exist.

Values 1 to 65535

queue-override

Syntax [**no**] **queue-override**

Context config>service>ies>if>sap>egress
config>service>ies>if>sap>ingress
config>service>vpls>sap>egress
config>service>vpls>sap>ingress

config>service>vprn>if>sap>egress
config>service>vprn>if>sap>ingress

Description This command enables the context to configure override values for the specified SAP egress QoS queue. These values override the corresponding ones specified in the associated SAP egress QoS policy.

queue

Syntax **queue** *queue-id* [**create**]
no queue *queue-id*

Context config>service>ies>if>sap>egress>queue-override
config>service>ies>if>sap>ingress>queue-override
config>service>vpls>sap>egress>queue-override
config>service>vpls>sap>ingress>queue-override
config>service>vpls>sap>egress>queue-override
config>service>vpls>sap>ingress>queue-override

Description This command specifies the ID of the queue whose parameters are to be overridden.

The **no** form of this command removes the queue ID from the configuration.

Parameters *queue-id* — Specifies the queue ID whose parameters are to be overridden.

Values 1 to 8

create — Keyword used to create the queue ID. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

adaptation-rule

Syntax **adaptation-rule** [**pir** *adaptation-rule*] [**cir** *adaptation-rule*]

Context config>service>ies>if>sap>egress>queue-override>queue
config>service>ies>if>sap>ingress>queue-override>queue
config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>sap>ingress>queue-override>queue

Description This command overrides specific attributes of the specified queue's adaptation rule parameters. The adaptation rule controls the method used by the system to derive the operational CIR and PIR settings when the queue is provisioned in hardware. For the CIR and PIR parameters individually, the system attempts to find the best operational rate depending on the defined constraint.

The **no** form of this command removes any explicitly defined constraints used to derive the operational CIR and PIR created by the application of the policy. When a specific **adaptation-rule** is removed, the default constraints for **rate** and **cir** apply.

Parameters **pir** — Specifies the constraints enforced when adapting the PIR rate defined within the **queue** *queue-id* **rate** command. The **pir** parameter requires a qualifier that defines the constraint used when deriving the operational PIR for the queue. When the **rate** command is not specified, the default applies.

cir — Specifies the constraints enforced when adapting the CIR rate defined within the **queue** *queue-id* **rate** command. The **cir** parameter requires a qualifier that defines the constraint used when deriving the operational CIR for the queue. When the **cir** parameter is not specified, the default constraint applies.

adaptation-rule — Specifies the CIR and PIR adaptation rules.

Values max — The **max** (maximum) option is mutually exclusive with the **min** and **closest** options. When **max** is defined, the operational PIR for the queue is equal to or less than the administrative rate specified using the **rate** command.

min — The **min** (minimum) option is mutually exclusive with the **max** and **closest** options. When **min** is defined, the operational PIR for the queue is equal to or greater than the administrative rate specified using the **rate** command.

closest — The **closest** parameter is mutually exclusive with the **min** and **max** parameter. When **closest** is defined, the operational PIR for the queue is the rate closest to the rate specified using the **rate** command.

avg-frame-overhead

Syntax **avg-frame-overhead** *percent*
no avg-frame-overhead

Context config>service>ies>if>sap>egress>queue-override>queue
config>service>ies>if>sap>ingress>queue-override>queue
config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>sap>ingress>queue-override>queue

Description This command configures the average frame overhead to define the average percentage that the offered load to a queue expands during the frame encapsulation process before sending traffic on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a Sonet or SDH port or channel. Queues operating on egress Ethernet ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for Inter-Frame Gap).

When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:

- Offered-Load — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet-based offered-load.

- Frame-encapsulation overhead — Using the avg-frame-overhead parameter, the frame-encapsulation overhead is simply the queue's current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10,000 octets and the avg-frame-overhead equals 10%, the frame-encapsulation overhead would be 10,000 x 0.1 or 1,000 octets.

For egress Ethernet queues, the frame-encapsulation overhead is calculated by multiplying the number of offered-packets for the queue by 20 bytes. If a queue was offered 50 packets then the frame-encapsulation overhead would be 50 x 20 or 1,000 octets.

- Frame-based offered-load — The frame-based offered-load is calculated by adding the offered-load to the frame-encapsulation overhead. If the offered-load is 10,000 octets and the encapsulation overhead was 1,000 octets, the frame-based offered-load would equal 11,000 octets.
- Packet to frame factor — The packet to frame factor is calculated by dividing the frame-encapsulation overhead by the queue's offered-load (packet-based). If the frame-encapsulation overhead is 1,000 octets and the offered-load is 10,000 octets then the packet to frame factor would be 1,000 / 10,000 or 0.1. When in use, the avg-frame-overhead is the same as the packet to frame factor making this calculation unnecessary.
- Frame-based CIR — The frame-based CIR is calculated by multiplying the packet to frame factor with the queue's configured CIR, then adding that result to that CIR. If the queue CIR is set at 500 octets and the packet to frame factor equals 0.1, the frame-based CIR would be 500 x 1.1 or 550 octets.
- Frame-based within-CIR offered-load — The frame-based within-CIR offered-load is the portion of the frame-based offered-load considered to be within the frame-based CIR. The frame-based within-CIR offered-load is the lesser of the frame-based offered-load and the frame-based CIR. If the frame-based offered-load equaled 11000 octets and the frame-based CIR equaled 550 octets, the frame-based within-CIR offered-load would be limited to 550 octets. If the frame-based offered-load equaled 450 octets and the frame-based CIR equaled 550 octets, the frame-based within-CIR offered-load would equal 450 octets (or the entire frame-based offered-load).

As a special case, when a queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame-based within-CIR offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-CIR pass.

- Frame-based PIR — The frame-based PIR is calculated by multiplying the packet to frame factor with the queue's-configured PIR, then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame factor equals 0.1, the frame-based PIR would be 7,500 x 1.1 or 8,250 octets.
- Frame-based within-pir offered-load — The frame-based within-pir offered-load is the portion of the frame-based offered-load considered to be within the frame-based PIR. The frame-based within-pir offered-load is the lesser of the frame-based offered-load and the frame-based PIR. If the frame-based offered-load equaled 11,000 octets and the frame-based PIR equaled 8250 octets, the frame-based within-pir offered-load would be limited to 8,250 octets. If the frame-based offered-load equaled 7,000 octets and the frame-based PIR equaled 8,250 octets, the frame-based within-pir offered load would equal 7,000 octets.

Port Scheduler Operation Using Frame Transformed Rates — The port scheduler uses the frame-based rates to figure the maximum rates that each queue may receive during the within-CIR and above-CIR bandwidth allocation passes. During the within-CIR pass, a queue may receive up to its frame-based within-CIR offered-load. The maximum it may receive during the above-CIR pass is the difference between the frame-based within-pir offered load and the amount of actual bandwidth allocated during the within-CIR pass.

SAP and Subscriber SLA-Profile Average Frame Overhead Override — The average frame overhead parameter on a sap-egress may be overridden on an individual egress queue basis; on each SAP and within the sla-profile policy used by subscribers. An avg-frame-overhead command may be defined under the queue-override context for each queue. When overridden, the queue instance uses its local value for the average frame overhead instead of the sap-egress defined overhead.

The **no** form of this command restores the average frame overhead parameter for the queue to the default value of 0 percent. When set to 0, the system uses the packet-based queue statistics for calculating port scheduler priority bandwidth allocation. If the **no avg-frame-overhead** command is executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

Default avg-frame-overhead 0

Parameters *percent* — This parameter sets the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues. This parameter only applies to the 7450 ESS and 7750 SR.

Values 0.00 to 100.00

cbs

Syntax **cbs** *size-in-kbytes*
no cbs

Context config>service>ies>if>sap>egress>queue-override>queue
config>service>ies>if>sap>ingress>queue-override>queue
config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>sap>ingress>queue-override>queue

Description This command overrides specific attributes of the specified queue's CBS parameters.

It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a given access port egress buffer pool. Oversubscription may be desirable due to the potential large number of service queues and the economy of statistical multiplexing the individual queue's CBS settings into the defined reserved total.

When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high and low priority RED slopes on the pool, causing them to miscalculate when to start randomly drop packets.

The **no** form of this command returns the CBS size to the default value.

Parameters *size-in-kbytes* — Specifies the size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10 kbytes is desired, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).

Values 0 to 1048576, default

drop-tail

Syntax **drop-tail**

Context config>service>ies>if>sap>egress>queue-override>queue
config>service>ies>if>sap>ingress>queue-override>queue
config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>sap>ingress>queue-override>queue

Description This command enables the context to configure queue drop tail parameters.

low

Syntax **low**

Context config>service>ies>if>sap>egress>queue-override>queue>drop-tail
config>service>ies>if>sap>ingress>queue-override>queue>drop-tail
config>service>vpls>sap>egress>queue-override>queue>drop-tail
config>service>vpls>sap>ingress>queue-override>queue>drop-tail

Description This command enables the context to configure the queue **low drop-tail** parameters. The low drop tail defines the queue depth beyond which out-of-profile packets are not accepted into the queue and are discarded.

percent-reduction-from-mbs

Syntax **percent-reduction-from-mbs** *percent*
no percent-reduction-from-mbs

Context config>service>ies>if>sap>egress>queue-override>queue>drop-tail>low

config>service>ies>if>sap>ingress>queue-override>queue>drop-tail>low
config>service>vpls>sap>egress>queue-override>queue>drop-tail>low
config>service>vpls>sap>ingress>queue-override>queue>drop-tail>low

Description This command overrides the low queue drop tail as a percentage reduction from the MBS of the queue. For example, if a queue has an MBS of 600 kbytes and this percentage is configured to be 30% for the low drop tail, then the low drop tail is at 420 kbytes and out-of-profile packets are not accepted into the queue if its depth is greater than this value, and discarded.

Parameters *percent* — Specifies the percentage reduction from the MBS for a queue drop tail.

Values 0 to 100, default

mbs

Syntax **mbs** *size* [**bytes** | **kilobytes**]
no mbs

Context config>service>ies>if>sap>egress>queue-override>queue
config>service>ies>if>sap>ingress>queue-override>queue
config>service>vprn>sap>egress>queue-override>queue
config>service>vprn>sap>ingress>queue-override>queue

Description This command overrides specific attributes of the specified queue's MBS parameters. The MBS is a mechanism to override the default maximum size for the queue.

The sum of the MBS for all queues on an egress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer is available when needed or that the packets RED slope is not forced the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

The **no** form of this command returns the MBS size assigned to the queue.

Default mbs default

Parameters *size* — This required parameter specifies that the MBS is expressed as an integer representing the required size in either bytes or kilobytes. The default is **kilobytes**. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly to define the size. By specifying the keyword **default** sets the MBS to its default value.

Values 0 to 1073741824, default

bytes — Specifies that the value given for *size* is interpreted as the queue's MBS value in bytes.

kilobytes — Specifies that the value given for *size* is interpreted as the queue's MBS value in kb/s.

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

Context config>service>ies>if>sap>egress>queue-override>queue
config>service>ies>if>sap>ingress>queue-override>queue
config>service>vpls>sap>egress>queue-override>queue
config>service>vpls>sap>ingress>queue-override>queue

Description This command overrides specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.

The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile, then out-of-profile, packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at any time, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of this command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (max, 0).

Default rate max cir 0

Parameters *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.

Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 to 2000000000

Default max

cir-rate — Overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.

Fractional values are not allowed and must be given as a positive integer. The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers or queues.

Values 0 to 2000000000, max, sum

Default 0

scheduler-override

Syntax [**no**] **scheduler-override**

Context config>service>vpls>sap>egress
config>service>vpls>sap>ingress

Description This command enables the context to configure the set of attributes whose values have been overridden via management on this virtual scheduler. Clearing a given flag returns the corresponding overridden attribute to the value defined on the SAP's ingress scheduler policy.

The **no** form of this command removes scheduler parameters from the configuration.

scheduler

Syntax **scheduler** *scheduler-name* [**create**]
no scheduler *scheduler-name*

Context config>service>vpls>sap>egress>sched-override

Description This command overrides specific attributes of the specified scheduler name.

A scheduler defines a bandwidth control that limits each child (other schedulers, policers and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created has policers, queues or other schedulers defined as child associations. The scheduler can be a child which takes bandwidth from a scheduler in a higher tier. A total of 32 schedulers can be created within a single scheduler policy with no restriction on the distribution between the tiers.

Each scheduler must have a unique name within the context of the scheduler policy; however the same name can be reused in multiple scheduler policies. If *scheduler-name* already exists within the policy tier level (regardless of the inclusion of the keyword **create**), the context changes to that scheduler name for the purpose of editing the scheduler parameters. Modifications made to an existing scheduler are executed on all instantiated schedulers created through association with the policy of the edited scheduler. This can cause policers, queues or schedulers to become orphaned (invalid parent association) and adversely affect the ability of the system to enforce service level agreements (SLAs).

If the *scheduler-name* exists within the policy on a different tier (regardless of the inclusion of the keyword create), an error occurs and the current CLI context does not change.

If the *scheduler-name* does not exist in this or another tier within the scheduler policy, it is assumed that an attempt is being made to create a scheduler of that name. The success of the command execution is dependent on the following:

Step 1. The maximum number of schedulers has not been configured.

Step 2. The provided *scheduler-name* is valid.

Step 3. The **create** keyword is entered with the command if the system is configured to require it (enabled in the **environment create** command).

When the maximum number of schedulers has been exceeded on the policy, a configuration error occurs and the command does not execute, nor does the CLI context change.

If the provided scheduler-name is invalid according to the criteria below, a name syntax error occurs, the command does not execute, and the CLI context does not change.

The **no** form of this command removes the scheduler name from the configuration.

Parameters *scheduler-name* — Specifies name of the scheduler

Values Valid names consist of any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

create — This optional keyword explicitly specifies that it is acceptable to create a scheduler with the given *scheduler-name*. If the **create** keyword is omitted, **scheduler-name** is not created when the system environment variable create is set to true. This safeguard is meant to avoid accidental creation of system objects (such as schedulers) while attempting to edit an object with a mistyped name or ID. The keyword has no effect when the object already exists.

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

Context config>service>vpls>sap>egress>queue-override>queue

config>service>vpls>sap>ingress>queue-override>queue

Description This command overrides specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters.

The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile, then out-of-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of this command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default rate max cir 0

Parameters *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.
Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 to 2000000000

Default **max**

cir-rate — Overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.
Fractional values are not allowed and must be given as a positive integer.

Values 0 to 2000000000, **max**

Default 0

scheduler-policy

Syntax **scheduler-policy** *scheduler-policy-name*
no scheduler-policy

Context config>service>vpls>sap>ingress
config>service>vpls>sap>egress
config>service>ies>if>sap>egress
config>service>ies>if>sap>ingress
config>service>vprn>if>sap>egress
config>service>vprn>if>sap>ingress

Description This command applies an existing scheduler policy to an ingress or egress scheduler used by SAP queues associated with this multi-service customer site. The schedulers defined in the scheduler policy can only be created once the customer site has been appropriately assigned to a chassis port, channel or slot. Scheduler policies are defined in the **config>qos>scheduler-policy** *scheduler-policy-name* context.

The **no** form of this command removes the configured ingress or egress scheduler policy from the multi-service customer site. When the policy is removed, the schedulers created due to the policy are removed also making them unavailable for the ingress SAP queues and egress SAP policers and queues associated with the customer site. Policers and queues that lose their parent scheduler association are deemed to be orphaned and are no longer subject to a virtual scheduler. The SAPs that have policers or queues reliant on the removed schedulers enter into an operational state depicting the orphaned status of one or more policers or queues. When the **no scheduler-policy** command is executed, the customer site's ingress or egress node will not contain an applied scheduler policy.

Parameters *scheduler-policy-name* — Specifies that the *scheduler-policy-name* is applied to an existing scheduler policy that was created in the **config>qos>scheduler-policy** *scheduler-policy-name* context to create the hierarchy of ingress or egress virtual schedulers. The scheduler names defined within the policy are created and made available to any ingress or egress queues created on associated SAPs.

Values Any existing valid scheduler policy name.

ingress

Syntax **ingress**

Context config>service>vpls>sap
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp
config>service>ies>if>sap

Description This command enables the context to configure ingress Quality of Service (QoS) policies and filter policies.

If no QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing. If no ingress filter is defined, no filtering is performed.

match-qinq-dot1p

Syntax **match-qinq-dot1p** {**top** | **bottom**}
no match-qinq-dot1p

Context config>service>vpls>sap>ingress
config>service>ies>if>sap>ingress

Description This command specifies which dot1Q tag position dot1P bits in a QinQ encapsulated packet should be used to evaluate dot1P QoS classification.

The **match-qinq-dot1p** command allows the top or bottom PBits to be used when evaluating the applied sap-ingress QoS policy's dot1P entries. The **top** and **bottom** keywords specify which position should be evaluated for QinQ encapsulated packets.

By default, the bottom-most service delineating dot1Q tag's dot1P bits are used. Table 8 defines the default behavior for dot1P evaluation when the **match-qinq-dot1p** command is not executed.

Table 8 **Default QinQ and TopQ SAP Dot1P Evaluation**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ/TopQ	TopQ	TopQ PBits
QinQ/TopQ	TopQ BottomQ	TopQ PBits
QinQ/TopQ	TopQ BottomQ	BottomQ PBits

The **no** form of this command restores the default dot1p evaluation behavior for the SAP.

Default no match-qinq-dot1p (no filtering based on p-bits)

(top or bottom must be specified to override the default QinQ dot1p behavior)

Parameters **top** — The top parameter is mutually exclusive to the bottom parameter. When the **top** parameter is specified, the topmost PBits are used (if existing) to match any dot1p dot1p-value entries. Table 9 defines the dot1p evaluation behavior when the **top** parameter is specified.

Table 9 **Top Position QinQ and TopQ SAP Dot1P Evaluation**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ/TopQ	TopQ	TopQ PBits
QinQ/TopQ	TopQ BottomQ	TopQ PBits
QinQ/QinQ	TopQ BottomQ	TopQ PBits

bottom — The **bottom** parameter is mutually exclusive to the **top** parameter. When the bottom parameter is specified, the bottom most PBits are used (if existing) to match any dot1p dot1p-value entries. Table 10 defines the dot1p evaluation behavior when the bottom parameter is specified.

Table 10 **Bottom Position QinQ and TopQ SAP Dot1P Evaluation**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	BottomQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None

Table 10 **Bottom Position QinQ and TopQ SAP Dot1P Evaluation**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ/TopQ	TopQ	TopQ PBits
QinQ/TopQ	TopQ BottomQ	BottomQ PBits
QinQ/QinQ	TopQ BottomQ	BottomQ PBits

qos

Syntax **qos** *policy-id* [**shared-queuing** | **multipoint-shared**] [**fp-redirect-group** *queue-group-name* **instance** *instance-id*]

qos *policy-id* [**shared-queuing** | **multipoint-shared**]

no qos

Context config>service>ies>if>sap>ingress
config>service>vprn>if>sap>ingress

Description This command associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP).

QoS ingress and egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP. If the policy ID does not exist, an error is returned.

The **qos** command is used to associate both ingress and egress QoS policies. The qos command only allows ingress policies to be associated on SAP ingress and egress policies on SAP egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type returns an error.

By default, no specific QoS policy is associated with the SAP for ingress or egress, so the default QoS policy is used.

The no form of this command removes the QoS policy association from the SAP, and the QoS policy reverts to the default.

The **no** form of this command removes the QoS policy association from the SAP or IP interface, and the QoS policy reverts to the default.

Parameters *policy-id* — The ingress/egress policy ID to associate with SAP or IP interface on ingress/egress. The policy ID must already exist.

1 to 65535

shared-queuing — Specifies the ingress shared queue policy used by this SAP. When the value of this object is null it means that the SAP uses individual ingress QoS queues instead of the shared ones.

multipoint-shared — This keyword specifies that this queue-id is for multipoint forwarded traffic only. This queue-id can only be explicitly mapped to the forwarding class multicast, broadcast, or unknown unicast ingress traffic. Attempting to map forwarding class unicast traffic to a multipoint queue generates an error; no changes are made to the current unicast traffic queue mapping.

A queue must be created as multipoint. The multipoint designator cannot be defined after the queue is created. If an attempt is made to modify the command to include the multipoint keyword, an error is generated and the command is not executed.

The multipoint keyword can be entered in the command line on a preexisting multipoint queue to edit queue ID parameters.

Default Present (the queue is created as non-multipoint).

Values **Multipoint** or not present.

queue-group-name — Specifies the name of the queue group template to be instantiated on the forwarding plane of the IOM/IMM/XMA, up to 32 characters. The queue-group-name must correspond to a valid ingress queue group template name, configured in the **config>qos>queue-group-templates** context.

instance-id — Specifies the instance of the named queue group to be created on the IOM/IMM/XMA ingress forwarding plane.

multi-service-site

Syntax **multi-service-site** *customer-site-name*
no multi-service-site

Context config>service>ies>sap
config>service>vpls>sap
config>service>vprn>if>sap

Description This command associates the SAP with a *customer-site-name*. If the specified *customer-site-name* does not exist in the context of the service customer ID an error occurs and the command is not executed. If *customer-site-name* exists, the current and future defined queues on the SAP (ingress and egress) attempts to use the scheduler hierarchies created within *customer-site-name* as parent schedulers.

This command is mutually exclusive with the SAP ingress and egress scheduler policy commands. If a scheduler policy has been applied to either the ingress or egress nodes on the SAP, the **multi-service-site** command fails without executing. The locally applied scheduler policies must be removed prior to executing the **multi-service-site** command.

The **no** form of this command removes the SAP from any multi-service customer site the SAP belongs to. Removing the site can cause existing or future policers and queues to enter an orphaned state.

Parameters *customer-site-name* — Specifies an existing customer site name, up to 32 characters. If the *customer-site-name* exists and local scheduler policies have not been applied to the SAP, the current and future policers queues defined on the SAP looks for their parent schedulers within the scheduler hierarchies defined in the customer-site-name.

static-host

Syntax **static-host ip** *ip-address* [**mac** *ieee-address*] [**create**]
static-host mac *ieee-address* [**create**]
no static-host ip *ip-address* **mac** *ieee-address*
no static-host all [**force**]
no static-host ip *ip-address*

Context config>service>ies>if>sap
config>service>vpls>sap
config>service>vprn>if>sap

Description This command creates a static subscriber host for the SAP. Static subscriber hosts may be used by the system for various purposes. Applications within the system that make use of static host entries include anti-spoof, ARP reply agent and source MAC population into the VPLS forwarding database.

Multiple static hosts may be defined on the SAP. Each host is identified by either a source IP address, a source MAC address or both a source IP and source MAC address. Every static host definition must have at least one address defined, IP or MAC.

Static hosts can exist on the SAP even with anti-spoof and ARP reply agent features disabled. When enabled, each feature has different requirements for static hosts.

The **no** form of this command removes a static entry from the system. The specified *ip-address* and *mac-address* must match the host's exact IP and MAC addresses as defined when it was created. When a static host is removed from the SAP, the corresponding anti-spoof filter entry and/or FDB entry is also removed.

Parameters *ip-address* — Specify this optional parameter when defining a static host. The IP address must be specified for **anti-spoof ip**, **anti-spoof ip-mac** and **arp-reply-agent** (**arp-reply-agent** is supported by the 7450 ESS only). Only one static host may be configured on the SAP with a given IP address.

mac-address — Specify this optional parameter when defining a static host. The MAC address must be specified for **anti-spoof mac**, and **anti-spoof ip-mac** and arp-reply-agent (a**rp-reply-agent** is supported by the 7450 ESS only). Multiple static hosts may be configured with the same MAC address given that each definition is distinguished by a unique IP address.

Every static host definition must have at least one address defined, IP or MAC.

force — Specifies the forced removal of the static host addresses.

create — Keyword used to create the static host instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

ancp-string

Syntax **ancp-string** *ancp-string*
no ancp-string

Context config>service>ies>if>sap>static-host
config>service>vpls>sap>static-host
config>service>vprn>if>sap>static-host

Description This command specifies the ANCP string associated to this SAP host.

The **no** form of this command reverts to the default.

Parameters *ancp-string* — Specifies the ANCP string up to 63 characters.

app-profile

Syntax **app-profile** *app-profile-name*
no app-profile

Context config>service>ies>if>sap>static-host
config>service>vpls>sap>static-host
config>service>vprn>if>sap

Description This command specifies an application profile name.

The **no** form of this command reverts to the default.

Parameters *app-profile-name* — Specifies the application profile name up to 32 characters.

inter-dest-id

Syntax **inter-dest-id** *intermediate-destination-id*
no inter-dest-id

Context config>service>ies>if>sap>static-host
config>service>vpls>sap>static-host
config>service>vprn>if>sap>static-host

Description This command specifies to which intermediate destination (for example a DSLAM) this host belongs.

The **no** form of this command reverts to the default.

Parameters *intermediate-destination-id* — Specifies the intermediate destination ID.

sla-profile

Syntax **sla-profile** *sla-profile-name*
no sla-profile

Context config>service>ies>if>sap>static-host
config>service>vpls>sap>static-host
config>service>vprn>if>sap>static-host

Description This command specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

The **no** form of this command reverts to the default.

Parameters *sla-profile-name* — Specifies the SLA profile name.

sub-profile

Syntax **sub-profile** *sub-profile-name*
no sub-profile

Context config>service>ies>if>sap>static-host
config>service>vpls>sap>static-host
config>service>vprn>if>sap>static-host

Description This command specifies an existing subscriber profile name to be associated with the static subscriber host.

The **no** form of this command reverts to the default.

Parameters *sub-profile-name* — Specifies the sub-profile name.

subscriber

Syntax **subscriber** *sub-ident*
no subscriber

Context config>service>ies>if>sap>static-host
config>service>vpls>sap>static-host
config>service>vprn>if>sap>static-host

Description This command specifies an existing subscriber identification profile to be associated with the static subscriber host.

Parameters *sub-ident* — Specifies the subscriber identification.

subscriber-sap-id

Syntax [**no**] **subscriber-sap-id**

Context config>service>ies>if>sap>static-host
config>service>vpls>sap>static-host
config>service>vprn>if>sap>static-host

Description This command enables using the SAP ID as subscriber id.

Parameters **subscriber-sap-id** — Specifies to use the sap-id as the subscriber-id.

egress

Syntax **egress**

Context config>service>ies>if>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command enables the context to configure egress SDP parameters.

vc-label

Syntax [**no**] **vc-label** *egress-vc-label*

Context config>service>ies>if>spoke-sdp>egress
config>service>vpls>mesh-sdp>egress
config>service>vpls>spoke-sdp>egress

Description This command configures egress VC label parameters.

The **no** form of this command reverts to the default.

Parameters *vc-label* — Specifies a VC egress value that indicates a specific connection.

Values 16 to 1048575

ingress

Syntax **ingress**

Context config>service>ies>if>spoke-sdp
config>service>vpls>mesh-sdp

config>service>vpls>spoke-sdp

Description This command enables the context to configure ingress SDP parameters.

vc-label

Syntax [**no**] **vc-label** *ingress-vc-label*

Context config>service>ies>if>spoke-sdp>ingress
config>service>vpls>mesh-sdp>ingress
config>service>vpls>spoke-sdp>ingress

Description This command configures the ingress VC label.

The **no** form of this command reverts to the default.

Parameters *vc-label* — Specifies a VC ingress value that indicates a specific connection.

Values 2048 to 18431

secondary

Syntax **secondary** {*ip-address/mask* | *ip-address netmask*} [**broadcast all-ones** | **host-ones**] [**igp-inhibit**]
no secondary *ip-address*

Context config>service>ies>if

Description This command assigns a secondary IP address or IP subnet/broadcast address format to the interface.

The **no** form of this command reverts to the default.

Parameters *ip-address* — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that is used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

mask — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that is used in a logical and function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252.

Note: A mask of 255.255.255.255 is reserved for system IP addresses.

netmask — Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains.

broadcast — Overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is **host-ones** which indicates a subnet broadcast address. Use this parameter to change the broadcast address to **all-ones** or revert back to a broadcast address of **host-ones**.

The broadcast format on an IP interface can be specified when the IP address is assigned or changed.

This parameter does not affect the type of broadcasts that can be received by the IP interface. A host sending either the local broadcast (**all-ones**) or the valid subnet broadcast address (**host-ones**) is received by the IP interface. (Default: host-ones)

all-ones — Specifies the broadcast address used by the IP interface for this IP address is 255.255.255.255, also known as the local broadcast.

host-ones — Specifies that the broadcast address used by the IP interface for this IP address is the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the *ip-address* and the *mask-length* or *mask* with all the host bits set to binary one. This is the default broadcast address used by an IP interface.
The **broadcast** parameter within the **address** command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the **broadcast** type to **host-ones** after being changed to **all-ones**, the **address** command must be executed with the **broadcast** parameter defined.

igp-inhibit — Signals that the given secondary IP interface should not be recognized as a local interface by the running IGP. For OSPF and IS-IS, this means that the specified secondary IP interfaces are not injected and used as passive interfaces and are not advertised as internal IP interfaces into the IGP's link state database. For RIP, this means that these secondary IP interfaces do not source RIP updates.

static-arp

Syntax **static-arp** *ip-address ieee-mac-address*
no static-arp *ip-address ieee-mac-address*

Context config>service>ies>if
config>service>vprn>if

Description This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. This static ARP appears in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface.

If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address is replaced with the new MAC address.

The **no** form of this command removes a static ARP entry.

Parameters *ip-address* — Specifies the IP address for the static ARP in IP address dotted decimal notation.

ieee-mac-address — Specifies the 48-bit MAC address for the static ARP in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

static-mac

Syntax **static-mac** *ieee-mac-address* [**create**]
no static-mac *ieee-mac-address*

Context config>service>vpls>sap
config>service>vpls>mesh-sdp
config>service>vpls>spoke-sdp

Description This command creates a remote static MAC entry in the Virtual Private LAN Service (VPLS) forwarding database (FDB) associated with the Service Distribution Point (SDP).

In a VPLS service, MAC addresses are associated with a Service Access Point (SAP) or with a Service Distribution Point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Remote static MAC entries create a permanent MAC address to SDP association in the forwarding database for the VPLS instance so that MAC address is not learned on the edge device.

Note: Static MAC definitions on one edge device are not propagated to other edge devices participating in the VPLS instance, that is, each edge device has an independent forwarding database for the VPLS.

Only one static MAC entry (local or remote) can be defined per MAC address per VPLS instance.

By default, no static MAC address entries are defined for the SDP.

The **no** form of this command deletes the static MAC entry with the specified MAC address associated with the SDP from the VPLS forwarding database.

Parameters *ieee-mac-address* — Specifies the 48-bit MAC address for the static ARP in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

create — Keyword used to create the static MACinstance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

vlan-vc-tag

Syntax **vlan-vc-tag** *vlan-id*
no vlan-vc-tag [*vlan-id*]

Context config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp

Description This command specifies an explicit dot1q value used when encapsulating to the SDP far end. When signaling is enabled between the near and far end, the configured dot1q tag can be overridden by a received TLV specifying the dot1q value expected by the far end. This signaled value must be stored as the remote signaled dot1q value for the binding. The provisioned local dot1q tag must be stored as the administrative dot1q value for the binding.

When the dot1q tag is not defined, the default value of zero is stored as the administrative dot1q value. Setting the value to zero is equivalent to not specifying the value.

The **no** form of this command disables the command.

Parameters *vlan-id* — Specifies a valid VLAN identifier to bind an 802.1Q VLAN tag ID.

mesh-sdp

Syntax **mesh-sdp** *sdp-id*[*:vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**]
mesh-sdp *sdp-id*[*:vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**] **leaf-ac**
mesh-sdp *sdp-id*[*:vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**] **root-leaf-tag**
no mesh-sdp *sdp-id*[*:vc-id*]

Context config>service>vpls

Description This command binds a VPLS service to an existing Service Distribution Point (SDP).

Mesh SDPs bound to a service are logically treated like a single bridge “port” for flooded traffic where flooded traffic received on any mesh SDP on the service is replicated to other “ports” (spoke SDPs and SAPs) and not transmitted on any mesh SDPs.

Note: This command creates a binding between a service and an SDP. The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service is down.

The SDP must already be defined in the **config>service>sdp** context in order to associate the SDP with an Epipe or VPLS service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end router devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

Special Cases **VPLS —** Several SDPs can be bound to a VPLS. Each SDP must be destined to a different router. If two *sdp-id* bindings terminate on the same router, an error occurs and the second SDP binding is rejected.

Parameters *sdp-id* — Specifies an SDP identifier.

Values 1 to 17407

vc-id — Specifies a virtual circuit identifier. This value is used to validate the VC ID portion of each mesh SDP binding defined in the service. The default value of this object is equal to the service ID. Any value may be used for the *vc-id* when there is no existing mesh SDP within the service; if a mesh SDP exists then all other mesh SDPs in the service must be configured with the same *vc-id*.

Values 1 to 4294967295

vc-type — Specifies to override the default VC type signaled for the spoke or mesh binding to the far end of the SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment. A change of the binding's VC type causes the binding to signal the new VC type to the far end when signaling is enabled.

VC types are derived according to IETF *draft-martini-l2circuit-trans-mpls*.

The VC type value for Ethernet is 0x0005.

The VC type value for an Ethernet VLAN is 0x0004.

ether — Defines the VC type as Ethernet. The **vlan** keyword is mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings. Defining Ethernet is the same as executing **no vc-type** and restores the default VC type for the spoke SDP binding. (hex 5)

vlan — Defines the VC type as VLAN. The **ether** keyword is mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings.

root-leaf-tag — Specifies a tagging mesh SDP under an E-Tree VPLS. When a tag SDP binding is required, it is created with a root-leaf-tag flag. Only VLAN tag SDP bindings are supported. The VLAN type must be set to VC VLAN type. The root-leaf-tag parameter indicates this SDP binding is a tag SDP that uses a default VID 1 for root and 2 for leaf. The SDP binding tags egress E-Tree traffic with root and leaf VIDs as appropriate. Root and leaf VIDs are only significant between peering VPLS but the values must be consistent on each end. On ingress a tag SDP binding removes the VID tag on the interface between VPLS in the same E-Tree service. The tag SDP receives root tagged traffic and marks the traffic with a root indication internally. This option is not available on BGP EVPN-enabled E-Tree services.

leaf-ac — Specifies an access (AC) mesh SDP binding under a E-Tree VPLS as a leaf access (AC) SDP. The default E-Tree SDP type is a root AC if *leaf-ac* or *root-leaf-tag* is not specified at SDP binding creation. This option is only available when the VPLS is designated as an E-Tree VPLS. BGP EVPN-enabled E-Tree VPLS services support the **leaf-ac** option.

create — Keyword used to create the mesh SDP. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **ipipe**}] [**create**]
no spoke-sdp *sdp-id*[**:***vc-id*]

Context config>service>ies>if
config>service>vprn>if

Description This command binds a service to an existing Service Distribution Point (SDP).

A spoke SDP is treated like the equivalent of a traditional bridge “port” where flooded traffic received on the spoke SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service is down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with an IES service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router. The spoke SDP must be shut down first before it can be deleted from the configuration.

Special Cases **IES —** At most, only one *sdp-id* can be bound to an IES service.

Parameters *sdp-id* — Specifies the SDP identifier. Allowed values are integers in the range of 1 and 17407 for existing SDPs.

vc-id — Specifies the virtual circuit identifier.

Values 1 to 4294967295

vc-type — Specifies the encapsulation and pseudowire type for the spoke SDP.

Values ether: specifies Ethernet pseudowire as the type of virtual circuit (VC) associated with the SDP binding

ipipe: specifies Ipipe pseudowire as the type of virtual circuit (VC) associated with the SDP binding

Default ether

create — Keyword used to create the spoke SDP. The create keyword requirement can be enabled or disabled in the **environment>create** context.

spoke-sdp

Syntax **spoke-sdp** *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*] [**create**] **endpoint** *endpoint* **root-leaf-tag**
spoke-sdp *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*] [**create**] **endpoint endpoint leaf-ac**
spoke-sdp *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*] [**create**] [**no-endpoint**]
spoke-sdp *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*] [**create**] **endpoint endpoint**
spoke-sdp *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*] [**create**] [**no-endpoint**] **leaf-ac**
spoke-sdp *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*] [**create**] [**no-endpoint**] **root-leaf-tag**
no spoke-sdp *sdp-id*[**:***vc-id*]

Context config>service>vpls

Description This command binds a service to an existing Service Distribution Point (SDP).

A spoke SDP is treated like the equivalent of a traditional bridge port where flooded traffic received on the spoke SDP is replicated on all other ports (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service is down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with a VPLS service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end 7450 ESS or 7750 SR devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

Special Cases **VPLS —** Several SDPs can be bound to a VPLS service. Each SDP must use unique *vc-ids*. An error message is generated if two SDP bindings with identical *vc-ids* terminate on the same router. Split horizon groups can only be created in the scope of a VPLS service.

Parameters *sdp-id* — Specifies the SDP identifier.

Values 1 to 17407

vc-id — Specifies the virtual circuit identifier.

Values 1 to 4294967295

vc-type — This command overrides the default VC type signaled for the spoke or mesh binding to the far end of the SDP. The VC type is a 15 bit-quantity containing a value which represents the type of VC. The actual signaling of the VC type depends on the signaling parameter defined for the SDP. If signaling is disabled, the **vc-type** command can still be used to define the dot1q value expected by the far-end provider equipment. A change of the binding's VC type causes the binding to signal the new VC type to the far end when signaling is enabled.

VC types are derived according to IETF *draft-martini-l2circuit-trans-mpls*.

The VC type value for Ethernet is 0x0005.

The VC type value for an Ethernet VLAN is 0x0004.

Values ether, vlan

ether — Defines the VC type as Ethernet. The **ethernet**, **vlan,** and **vpls** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings. Defining Ethernet is the same as executing **no vc-type** and restores the default VC type for the spoke SDP binding (hex 5)

vlan — Defines the VC type as VLAN. The **ethernet**, **vlan,** and **vpls** keywords are mutually exclusive. When the VC type is not defined then the default is Ethernet for spoke SDP bindings.
The VLAN VC-type requires at least one dot1Q tag within each encapsulated Ethernet packet transmitted to the far end.

group-name — Specifies the name of the split horizon group to which the SDP belongs.

root-leaf-tag — Specifies a SAP as a root leaf tag SDP. This option is only available when the VPLS is designated as an E-Tree VPLS.

leaf-tag-vid — Specifies to replace the outer SDP ID for leaf traffic. The leaf tag VID is only significant between peering VPLS but the values must be consistent on each end.

leaf-ac — Specifies a SDP as a leaf access (AC) SDP. The default E-Tree SAP type is root AC if **leaf-ac** (or **root-leaf-tag)** is not specified at SDP creation. This option is only available when the VPLS is designated as an E-Tree VPLS.

create — Keyword used to create the spoke SDP. The create keyword requirement can be enabled or disabled in the **environment>create** context.

block-on-mesh-failure

Syntax [**no**] **block-on-mesh-failure**

Context config>service>vpls>spoke-sdp

Description This command enables blocking (bring the spoke SDP to an operationally down state) after all configured mesh SDPs are in operationally down state. This event is signaled to corresponding T-LDP peer by withdrawing service label (status-bit-signaling non-capable peer) or by setting the PW not forwarding status bit in T-LDP message (status-bit-signaling capable peer).

The **no** form of this command reverts to the default.

filter

Syntax **filter ip** *ip-filter-id*
filter ipv6 *ipv6-filter-id*
filter mac *mac-filter-id*
no filter [**ip** *ip-filter-id*] [**mac** *mac-filter-id*] [**ipv6** *ipv6-filter-id*]

Context config>service>vpls>mesh-sdp>egress
config>service>vpls>mesh-sdp>ingress
config>service>vpls>spoke-sdp>egress
config>service>vpls>spoke-sdp>ingress

Description This command associates an IP filter policy or MAC filter policy with an ingress or egress SDP.

Filter policies control the forwarding and dropping of packets based on IP or MAC matching criteria. There are two types of filter policies: IP and MAC. Only one type may be applied to a SDP at a time.

The **filter** command is used to associate a filter policy with a specified filter ID with an ingress or egress SDP. The filter ID must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation fails and an error message returned.

In general, filters applied to SDP (ingress or egress) apply to all packets on the SDP. One exception is non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SDP. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

Special Cases **VPLS —** Both MAC and IP filters are supported on a VPLS service SDP.

Parameters **ip** *ip-filter-id* — Specifies IP filter policy. The filter ID must already exist within the created IP filters.

Values 1 to 65535

ipv6 *ipv6-filter-id* — Specifies the IPv6 filter policy. The filter ID must already exist within the created IPv6 filters.

Values 1 to 65535

mac *mac-filter-id* — Specifies the MAC filter policy. The specified filter ID must already exist within the created MAC filters. The filter policy must already exist within the created MAC filters.

Values 1 to 65535

interface

Syntax **interface** *ip-int-name* [**create**]
interface *ip-int-name* [**create**] **tunnel**
no interface *ip-int-name*

Context config>service>vprn

Description This command creates a logical IP routing interface for a Virtual Private Routed Network (VPRN). Once created, attributes like an IP address and service access point (SAP) can be associated with the IP interface.

The **interface** command, under the context of services, is used to create and maintain IP routing interfaces within VPRN service IDs. The **interface** command can be executed in the context of an VPRN service ID. The IP interface created is associated with the service core network routing instance and default routing table. The typical use for IP interfaces created in this manner is for subscriber internet access.

Interface names are case sensitive and must be unique within the group of defined IP interfaces defined for **config router interface** and **config service vprn interface**. Interface names must not be in the dotted decimal notation of an IP address. For example, the name “1.1.1.1” is not allowed, but “int-1.1.1.1” is allowed. Show commands for router interfaces use either interface names or the IP addresses. Use unique IP address values and IP address names to maintain clarity. It could be unclear to the user if the same IP address and IP address name values are used. Although not recommended, duplicate interface names can exist in different router instances.

The available IP address space for local subnets and routes is controlled with the **config router service-prefix** command. The **service-prefix** command administers the allowed subnets that can be defined on service IP interfaces. It also controls the prefixes that may be learned or statically defined with the service IP interface as the egress interface. This allows segmenting the IP address space into **config router** and **config service** domains.

When a new name is entered, a new logical router interface is created. When an existing interface name is entered, the user enters the router interface context for editing and configuration.

By default, there are no default IP interface names defined within the system. All VPRN IP interfaces must be explicitly defined. Interfaces are created in an enabled state.

The no form of this command removes IP the interface and all the associated configuration. The interface must be administratively shutdown before issuing the **no interface** command.

For VPRN services, the IP interface must be shutdown before the SAP on that interface may be removed. VPRN services do not have the **shutdown** command in the SAP CLI context. VPRN service SAPs rely on the interface status to enable and disable them.

Parameters *ip-int-name* — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed between double quotes.

Values 1 to 32 characters maximum

tunnel — Specifies that the interface is configured as tunnel interface, which could be used to terminate IPSec or GRE runnels in the private service.

create — Creates the interface. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

diameter-auth-policy

Syntax **diameter-auth-policy** *name*
no diameter-auth-policy

Context config>service>vpls>sap

Description This command is used to configure the Diameter NASREQ application policy to use for authentication.

The **no** form of this command reverts to the default value.

Parameters *name* — Specifies the name of the Diameter NASREQ application policy to use for authentication up to 32 characters.

sap

Syntax [**no**] **sap** *sap-id*

Context config>service>vprn>if

Description This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the router. Each SAP must be unique.

All SAPs must be explicitly created within a service or on an IP interface.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **config interface** *port-type port-id* **mode access** command. Channelized TDM ports are always access ports.

If a port is shut down, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service is discarded. The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP are also e deleted.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

port-id — Specifies the physical port ID in the *slot/mda/port* format.

If the card in the slot has Media Dependent Adapters (MDAs) installed, the *port-id* must be in the slot_number/MDA_number/port_number format. For example, 6/2/3 specifies port 3 on MDA 2 in slot 6.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels the port ID must include the channel ID. A period "." separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

create — Keyword used to create the SAP. The create keyword requirement can be enabled or disabled in the **environment>create** context.

limit-mac-move

Syntax **limit-mac-move** [**blockable** | **non-blockable**]
no limit-mac-move

Context config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command indicates whether or not the mac-move agent, when enabled using **config>service>vpls>mac-move** or **config>service>epipe>mac-move**, limits the MAC re-learn (move) rate on this SAP.

Default limit-mac-move blockable

Parameters **blockable** — Specifies that the agent monitors the MAC re-learn rate on the SAP, and it blocks it when the re-learn rate is exceeded.

non-blockable — Specifies that this SAP is not blocked, and another blockable SAP is blocked instead.

mac-pinning

Syntax [**no**] **mac-pinning**

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp

Description This command disables re-learning of MAC addresses on other SAPs within the VPLS. The MAC address remains attached to a given SAP for the duration of its age-timer.

The age of the MAC address entry in the FDB is set by the age timer. If **mac-aging** is disabled on a given VPLS service, any MAC address learned on a SAP/SDP with **mac-pinning** enabled remains in the FDB on this SAP or SDP forever.

Every event that would otherwise result in re-learning is logged (MAC address, original-SAP, new-SAP).

Note: MAC addresses learned during DHCP address assignment (DHCP snooping enabled) are not impacted by this command. MAC-pinning for such addresses is implicit.

The **no** form of this command enables re-learning.

Default When a SAP or spoke SDP is part of a Residential Split Horizon Group (RSHG), MAC pinning is activated at creation of the SAP. Otherwise MAC pinning is not enabled by default.

managed-vlan-list

Syntax **managed-vlan-list**

Context config>service>vpls>sap

Description This command enables the context to configure VLAN ranges to be managed by a management VPLS. The list indicates, for each SAP, the ranges of associated VLANs that are affected when the SAP changes state.

This command is only valid when the VPLS in which it is entered was created as a management VPLS.

range

Syntax [**no**] **range** *vlan-range*

Context config>service>vpls>sap>managed-vlan-list

Description This command configures a range of VLANs on an access port that are to be managed by an existing management VPLS.

This command is only valid when the VPLS in which it is entered was created as a management VPLS, and when the SAP in which it was entered was created on an Ethernet port with encapsulation type of dot1q or qinq, or on a SONET/SDH port with encapsulation type of bcp-dot1q.

To modify the range of VLANs, first the new range should be entered and afterwards the old range removed.

The **no** form of this command removes the VLAN range from this configuration.

Parameters *vlan-range* — Specifies the VLAN start value and VLAN end value. The end-vlan must be greater than start-vlan. The format is <start-vlan>-<end-vlan>.

Values start-vlan: 1 to 4094
end-vlan: 1 to 4094

track-srrp

Syntax **track-srrp** *srrp-id*
no track-srrp

Context config>service>vpls>sap

Description This is a capture SAP level command. This command is important in PPPoE deployments with MSAPs. PPPoE operation requires that the MAC address learned by the client at the very beginning of the session negotiation phase remains unchanged for the lifetime of the session (RFC 2516). This command ensures that the virtual MAC address used during the PPPoE session negotiation phase on the capture SAP is the same virtual MAC address that is used by the SRRP on the group interface on which the session is established. Therefore, it is mandated that the SRRP instance (and implicitly the group-interface) where the session belongs to is known in advance. If the group interface name for the session is returned by the RADIUS, it must be ensured that this group interface is the one on which the tracked SRRP instance is configured. PPPoE sessions on the same capture SAP cannot be shared across multiple group interfaces, but instead they all must belong to a single group interface that is known in advance.

The same restrictions apply to IPoE clients in MC Redundancy scenario if they are to be supported concurrently on the same capture SAP as PPPoE.

The supported capture SAP syntax is this:

sap <port-id>:X.* capture-sap

The capture SAP syntax that is not supported is this:

sap <port-id>:*.* capture-sap

The **no** form of this command removes the SRRP ID from this configuration.

Parameters *srrp-id* — Specifies the SRRP instance number.

Values 1 to 4294967295

mbs

Syntax **mbs** {*size-in-kbytes* | **default**}
no mbs

Context config>service>vpls>sap>ingress>queue-override>queue
config>service>vpls>sap>egress>queue-override>queue

Description This command overrides specific attributes of the specified queue's MBS parameters. The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueuing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer be available when needed or that the packets RED slope is not forced the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

The **no** form of this command returns the MBS size assigned to the queue to the default value.

Default mbs default

Parameters *size-in-kbytes* — Specifies the maximum number of kilobytes of buffering allowed for the queue. For a value of 100 kb/s, enter the value 100. A value of 0 causes the queue to discard all packets.

Values 0 to 1073741824 or default

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

Context config>service>vpls>sap>egress>sched-override>scheduler

Description This command overrides specific attributes of the specified scheduler rate. The **rate** command defines the maximum bandwidth that the scheduler can offer its policers, child queues or schedulers. The maximum rate is limited to the amount of bandwidth the scheduler can receive from its parent scheduler. If the scheduler has no parent, the maximum rate is assumed to be the amount available to the scheduler. When a parent is associated with the scheduler, the CIR parameter provides the amount of bandwidth to be considered during the parent scheduler's 'within CIR' distribution phase.

The actual operating rate of the scheduler is limited by bandwidth constraints other than its maximum rate. The scheduler's parent scheduler may not have the available bandwidth to meet the scheduler's needs or the bandwidth available to the parent scheduler could be allocated to other child schedulers or child queues on the parent based on higher priority. The children of the scheduler may not need the maximum rate available to the scheduler due to insufficient offered load or limits to their own maximum rates.

When a scheduler is defined without specifying a rate, the default rate is **max**. If the scheduler is a root scheduler (no parent defined), the default maximum rate must be changed to an explicit value. Without this explicit value, the scheduler assumes that an infinite amount of bandwidth is available and allow all child queues and schedulers to operate at their maximum rates.

The **no** form of this command returns the scheduler's PIR and CIR parameters to the value configured in the applied scheduler policy.

Default rate max cir sum

Parameters *pir-rate* — Specifies the PIR rates. The **pir** parameter accepts a step multiplier value that specifies the multiplier used to determine the PIR rate at which the queue operates. A value of 0 to 100000000 or the keyword **max** is accepted. Any other value results in an error without modifying the current PIR rate.

To calculate the actual PIR rate, the rate described by the queue's **rate** is multiplied by the *pir-rate*.

The SAP ingress context for PIR is independent of the defined forwarding class (fc) for the queue. The default **pir** and definable range is identical for each class. The PIR in effect for a queue defines the maximum rate at which the queue is allowed to forward packets in a given second, thus shaping the queue's output.

The PIR parameter for SAP ingress queues do not have a negate (**no**) function. To return the queue's PIR rate to the default value, that value must be specified as the PIR value.

Values 1 to 100000000, max

Default max

cir-rate — Specifies the CIR rate. The **cir** parameter accepts a step-multiplier value that specifies the multiplier used to determine the CIR rate at which the queue operate. A value of 0 to 250 or the keyword max is accepted. Any other value results in an error without modifying the current CIR rate.

To calculate the actual CIR rate, the rate described by the **rate pir** *pir-rate* is multiplied by the *cir-rate*. If the **cir** is set to max, then the CIR rate is set to infinity.

The context for CIR is dependent on the defined forwarding class (fc) for the queue. The default CIR and definable range is different for each class. The CIR in effect for a policer or queue defines both its profile (in or out) marking level as well as the relative importance compared to other queues for scheduling purposes during congestion periods.

Values 0 to 10000000, max, sum

Default sum

l2pt-termination

Syntax [**no**] **l2pt-termination**

Context config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This commands enables L2PT termination on a given SAP or spoke SDP. L2PT termination is supported only for STP BPDUs. PDUs of other protocols are discarded.

This feature can be enabled only if STP is disabled in the context of the given VPLS service.

The **no** form of this command disables the L2PT termination.

max-nbr-mac-addr

Syntax **max-nbr-mac-addr** *table-size*
no max-nbr-mac-addr

Context config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command specifies the maximum number of FDB entries for both learned and static MAC addresses for this SAP or spoke SDP.

When the configured limit has been reached, and discard-unknown-source has been enabled for this SAP or spoke SDP (see the **discard-unknown-source** command), packets with unknown source MAC addresses are discarded.

The **no** form of this command restores the global MAC learning limitations for the SAP or spoke SDP.

Parameters *table-size* — Specifies the maximum number of MAC entries in the FDB from this SAP.

Values 1 to 511999 for the 7750 SR
1 to 131071 for the 7450 ESS

discard-unknown-source

Syntax [**no**] **discard-unknown-source**

Context config>service>vpls>sap
config>service>vpls>spoke-sdp

Description When this command is enabled, packets received on a SAP or on a spoke SDP with an unknown source MAC address are dropped only if the maximum number of MAC addresses for that SAP or spoke SDP (see max-nbr-mac-addr) has been reached. If max-nbr-mac-addr has not been set for the SAP or spoke SDP, enabling discard-unknown-source has no effect.

When disabled, the packets are forwarded based on the destination MAC addresses.

The **no** form of this command causes packets with an unknown source MAC addresses to be forwarded by destination MAC addresses in VPLS.

2.11.2.2.1 VPLS Triple Play Commands

disable-aging

Syntax [**no**] **disable-aging**

Context config>service>vpls
config>service>vpls>spoke-sdp
config>service>vpls>sap

Description This command disables MAC address aging across a VPLS service or on a VPLS service SAP or spoke SDP.

Like in a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the VPLS forwarding database (FDB). The **disable-aging** command turns off aging for local and remote learned MAC addresses.

When **no disable-aging** is specified for a VPLS, it is possible to disable aging for specific SAPs and/or spoke SDPs by entering the **disable-aging command at the appropriate level.**

When the **disable-aging** command is entered at the VPLS level, the **disable-aging** state of individual SAPs or SDPs is ignored.

The **no** form of this command enables aging on the VPLS service.

disable-learning

Syntax [**no**] **disable-learning**

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command enables learning of new MAC addresses in the VPLS forwarding database (FDB) for the service instance, SAP instance or spoke SDP instance.

When **disable-learning** is enabled, a new source MAC addresses is not entered in the VPLS service forwarding database for both local and remote MAC addresses.

When **no disable-learning** is specified for a VPLS on a 7450 ESS, it is possible to disable learning for specific SAPs and/or spoke SDPs by entering the **disable-learning command at the appropriate level.**

When disabled, new source MAC addresses are learned and entered into the VPLS forwarding database. When the **no disable-learning** command is entered on VPLS level on a 7450 ESS, the **disable-learning** state of individual SAPs or spoke SDPs is ignored.

This parameter is mainly used in conjunction with the **discard-unknown** command.

The **no** form of this command enables learning of MAC addresses meaning that normal MAC learning is enabled.

discard-unknown

Syntax [**no**] **discard-unknown**

Context config>service>vpls

Description By default, packets with unknown destination MAC addresses are flooded. If **discard-unknown** is enabled at the VPLS level, packets with unknown destination MAC address are dropped instead (even when configured FDB size limits for VPLS or SAP are not yet reached).

The **no** form of this command allows flooding of packets with unknown destination MAC addresses in the VPLS.

fdb-table-high-wmark

Syntax [**no**] **fdb-table-high-wmark** *high-water-mark*

Context config>service>vpls
config>service>template>vpls-template

Description This command specifies the value to send logs and traps when the threshold is reached.

Parameters *high-water-mark* — Specifies the value to send logs and traps when the threshold is reached

Values 0 to 100

Default 95%

fdb-table-low-wmark

Syntax [**no**] **fdb-table-low-wmark** *low-water-mark*

Context config>service>vpls
config>service>template>vpls-template

Description This command specifies the value to send logs and traps when the threshold is reached.

Parameters *low-water-mark* — Specifies the value to send logs and traps when the threshold is reached

Values 0 to 100

Default 90%

fdb-table-size

Syntax **fdb-table-size** *table-size*
no fdb-table-size [*table-size*]

Context config>service>vpls
config>service>template>vpls-template

Description This command specifies the maximum number of MAC entries in the forwarding database (FDB) for the VPLS instance on this node.

The **fdb-table-size** specifies the maximum number of forwarding database entries for both learned and static MAC addresses for the VPLS instance.

The **no** form of this command returns the maximum FDB table size to default.

Default fdb-table-size 250

Parameters *table-size* — Specifies the number of entries permitted in the forwarding database for this VPLS instance

Values 7450 ESS, 7950 XRS, or 7750 SR-7, SR-12, SR-12e: 1 to 511999
7750 SR-e, SR-a: 1 to 250000

local-age

Syntax **local-age** *seconds*
no local-age

Context config>service>vpls

Description Specifies the aging time for locally learned MAC addresses in the forwarding database (FDB) for the Virtual Private LAN Service (VPLS) instance.

In a VPLS service, MAC addresses are associated with a Service Access Point (SAP) or with a Service Distribution Point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Like in a Layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the FDB. The **local-age** timer specifies the aging time for local learned MAC addresses.

The **no** form of this command returns the local aging timer to the default value.

Default local-age 300

Parameters *seconds* — Specifies the aging time for local MACs expressed in seconds.

Values 60 to 86400

mac-protect

Syntax **mac-protect**

Context config>service>vpls

Description This command enables the context to configure MAC protection. When enabled, the agent protects the MAC from being learned or re-learned on a SAP that has restricted learning enabled.

mac

Syntax [**no**] **mac** *ieee-address*

Context config>service>vpls>mac-protect

Description This command specifies the 48-bit IEEE 802.3 MAC address.

The **no** form of this command removes the IEEE address from the configuration.

Parameters *ieee-address* — Specifies the 48-bit MAC address in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers.

mac-subnet-length

Syntax **mac-subnet-length** *subnet-length*
no mac-subnet-length

Context config>service>vpls

Description This command specifies the number of bits to be considered when performing MAC learning (MAC source) and MAC switching (MAC destination). Specifically, this value identifies how many bits, starting from the beginning of the MAC address are used. For example, if the mask-value of 28 is used, MAC learning only performs a lookup for the first 28 bits of the source MAC address when comparing with existing FDB entries. Then, it installs the first 28 bits in the FDB while zeroing out the last 20 bits of the MAC address. When performing switching in the reverse direction, only the first 28 bits of the destination MAC address are used to perform a FDB lookup to determine the next hop.

The **no** form of this command switches back to full MAC lookup.

Default mac-subnet-length 48

Parameters *subnet-length* — Specifies the number of bits to be considered when performing MAC learning or MAC switching.

Values 24 to 48

mcr-default-gtw

Syntax **mcr-default-gtw**

Context config>service>vpls

Description This command enables the context to configure the default gateway information when using Dual Homing in L2-TPSDA. The IP and MAC address of the default gateway used for subscribers on an L2 MC-Ring are configured in this context. After a ring heals or fails, the system sends out a gratuitous ARP on an active ring SAP in order to attract traffic from subscribers on the ring with connectivity to that SAP.

ip

Syntax **ip** *ip-address*
no ip

Context config>service>vpls>mcr-default-gtw

Description This command relates to a system configured for Dual Homing in L2-TPSDA. It defines the IP address used when the system sends out a gratuitous ARP on an active SAP after a ring heals or fails in order to attract traffic from subscribers on the ring with connectivity to that SAP.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the IP address in a.b.c.d. format.

mac

Syntax **mac** *ieee-address*
no mac

Context config>service>vpls>mcr-default-gtw

Description This command relates to a system configured for Dual Homing in L2-TPSDA and defines the MAC address used when the system sends out a gratuitous ARP on an active SAP after a ring heals or fails in order to attract traffic from subscribers on the ring with connectivity to that SAP.

The **no** form of this command reverts to the default.

Parameters *ieee-address* — Specifies the address in xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx format (cannot be all zeros).

remote-age

Syntax **remote-age** *seconds*
no remote-age

Context config>service>vpls

Description This command specifies the aging time for remotely learned MAC addresses in the forwarding database (FDB) for the Virtual Private LAN Service (VPLS) instance.

In a VPLS service, MAC addresses are associated with a Service Access Point (SAP) or with a Service Distribution Point (SDP). MACs associated with a SAP are classified as local MACs, and MACs associated with an SDP are remote MACs.

Like in a layer 2 switch, learned MACs can be aged out if no packets are sourced from the MAC address for a period of time (the aging time). In each VPLS service instance, there are independent aging timers for local learned MAC and remote learned MAC entries in the FDB. The **remote-age** timer specifies the aging time for remote learned MAC addresses. To reduce the amount of signaling required between switches configure this timer larger than the **local-age** timer.

The **no** form of this command returns the remote aging timer to the default value.

Default remote-age 900

Parameters *seconds* — Specifies the aging time for remote MACs expressed in seconds.

Values 60 to 86400

service-mtu

Syntax **service-mtu** *octets*
no service-mtu

Context config>service>vpls

Description This command configures the service payload (Maximum Transmission Unit – MTU), in bytes, for the service. This MTU value overrides the service-type default MTU.

The **service-mtu** defines the payload capabilities of the service. It is used by the system to validate the SAP and SDP binding's operational state within the service.

The service MTU and a SAP's service delineation encapsulation overhead (for example, 4 bytes for a Dot1q tag) is used to derive the required MTU of the physical port or channel on which the SAP was created. If the required payload is larger than the port or channel MTU, then the SAP is placed in an inoperative state. If the required MTU is equal to or less than the port or channel MTU, the SAP is able to transition to the operative state.

When binding an SDP to a service, the service MTU is compared to the path MTU associated with the SDP. The path MTU can be administratively defined in the context of the SDP. The default or administrative path MTU can be dynamically reduced due to the MTU capabilities discovered by the tunneling mechanism of the SDP or the egress interface MTU capabilities based on the next hop in the tunnel path. If the service MTU is larger than the path MTU, the SDP binding for the service is placed in an inoperative state. If the service MTU is equal to or less than the path MTU, then the SDP binding is placed in an operational state.

In the event that a service MTU, port or channel MTU, or path MTU is dynamically or administratively modified, then all associated SAP and SDP binding operational states are automatically re-evaluated.

The **no** form of this command returns the default **service-mtu** for the indicated service type to the default value.

Default service-mtu 1514

The following table displays MTU values for specific VC types.

Table 11 **MTU values for specific VC types**

VC-Type	Example Service MTU	Advertised MTU
Ethernet	1514	1500
Ethernet (with preserved dot1q)	1518	1504
VPLS	1514	1500
VPLS (with preserved dot1q)	1518	1504

Table 11 **MTU values for specific VC types (Continued)**

VC-Type	Example Service MTU	Advertised MTU
VLAN (dot1p transparent to MTU value)	1514	1500
VLAN (QinQ with preserved bottom Qtag)	1518	1504

Parameters *octets* — Specifies the size of the MTU in octets, expressed as a decimal integer.

Values 1 to 9194

split-horizon-group

Syntax **split-horizon-group** [*group-name*] [*residential-group*]
no split-horizon-group

Context config>service>vpls

Description This command creates a new split horizon group for the VPLS instance. Traffic arriving on a SAP or spoke SDP within this split horizon group is not copied to other SAPs or spoke SDPs in the same split horizon group.

A split horizon group must be created before SAPs and spoke SDPs can be assigned to the group.

The split horizon group is defined within the context of a single VPLS. The same group name can be re-used in different VPLS instances.

Up to 30 split horizon groups can be defined per VPLS instance.

The **no** form of this command removes the group name from the configuration.

Parameters *group-name* — Specifies the name of the split horizon group to which the SDP belongs.

residential-group — Defines a split horizon group as a residential split horizon group (RSHG). Doing so entails that:

a. SAPs that are members of this Residential Split Horizon Group have:
   - Double-pass queuing at ingress as default setting (can be disabled)
   - STP disabled (can not be enabled)
   - ARP reply agent enabled by default (can be disabled)
   - MAC pinning enabled by default (can be disabled)
   - broadcast packets are discarded, blocking unknown, flooded traffic

b. Spoke SDPs that are members of this RSHG have:
   - Downstream multicast traffic supported
   - Double-pass queuing is not applicable

- STP is disabled (can be enabled)
- ARP reply agent is not applicable (dhcp-lease-states are not supported on spoke SDPs)
- MAC pinning enabled per default (can be disabled)

sap

Syntax **sap** *sap-id* [**split-horizon-group** *group-name*] [**create**] [**capture-sap**] [**eth-ring** *ring-index*]
sap *sap-id* [**split-horizon-group** *group-name*] [**create**] [**capture-sap**] [**eth-ring** *ring-index*] **leaf-ac**
sap *sap-id* [**split-horizon-group** *group-name*] [**create**] [**capture-sap**] [**eth-ring** *ring-index*] **root-leaf-tag leaf-tag** *leaf-tag*
no sap *sap-id*

Context config>service>vpls

Description This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the 7450 ESS or 7750 SR. Each SAP must be unique. All SAPs must be explicitly created within a service or on an IP interface.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **config interface** *port-type port-id* **mode access** command. Channelized TDM ports are always access ports (TDM applies to the 7750 SR only).

If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service is discarded. The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP are also deleted. For Internet Ethernet Service (IES), the IP interface must be shutdown before the SAP on that interface may be removed.

Special Cases **VPLS SAP —** A VPLS SAP can be defined with Ethernet ports, SONET/SDH or TDM channels (TDM is supported on the 7750 SR only).

A default SAP has the following format: port-id:*. This type of SAP is supported only on Ethernet MDAs and its creation is allowed only in the scope of Layer 2 services (Epipe and VPLS). This type of SAP is mutually exclusive with a SAP defined by explicit null encapsulation (for example, 1/1/1:0).

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

port-id — Specifies the physical port ID in the *slot/mda/port* format.

If the card in the slot has Media Dependent Adapters (MDAs) installed, the *port-id* must be in the slot_number/MDA_number/port_number [.channel] format. For example 6/2/3 specifies port 3 on MDA 2 in slot 6.

The *port-id* must reference a valid port type. When the *port-id* parameter represents SONET/SDH and TDM channels, the port ID must include the channel ID. A period "." separates the physical port from the *channel-id*. The port must be configured as an access port.

If the SONET/SDH port is configured as clear-channel then only the port is specified.

group-name — Specifies the name of the split horizon group to which the SAP belongs. This parameter applies to the 7450 ESS or 7750 SR only.

capture-sap — Specifies a capturing SAP in which triggering packets are sent to the CPM. Non-triggering packets captured by the capture SAP are dropped. This parameter applies to the 7450 ESS or 7750 SR only.

create — Keyword used to create a SAP instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

root-leaf-tag — Specifies a SAP as a root leaf tag SAP. Only SAPs of the form dot1q (for example, 1/1/1:X) or qinq (for example, 1/1/1:X.Y, 1/1/1:X.*) are supported. The default E-Tree SAP type is a root AC, if *root-leaf-tag* **(**or *leaf-ac***)** is not specified at SAP creation. This option is only available when the VPLS is designated as an E-Tree VPLS.

leaf-tag-vid — Specifies to replace the outer SAP-ID for leaf traffic. The leaf tag VID is only significant between peering VPLS but the values must be consistent on each end.

leaf-ac — Specifies a SAP as a leaf access (AC) SAP. The default E-Tree SAP type is root AC if **leaf-ac** (or **root-leaf-tag)** is not specified at SAP creation. This option is only available when the VPLS is designated as an E-Tree VPLS.

arp-host

Syntax **arp-host**

Context config>service>vpls>sap

Description This command enables the context to configure ARP host parameters.

The **no** form of this command reverts to the default value.

host-limit

Syntax **host-limit** *max-num-hosts*
no host-limit

Context config>service>vpls>sap>arp-host

Description This command configures the maximum number of ARP hosts.

The **no** form of this command returns the value to the default.

Default host-limit 1

Parameters *max-num-hosts* — Specifies the maximum number of ARP hosts allowed on this SAP.

Note: The operational maximum value may be smaller due to equipped hardware dependencies.

Values 1 to 131071

min-auth-interval

Syntax **min-auth-interval** *min-auth-interval*
no min-auth-interval

Context config>service>vpls>sap>arp-host

Description This command configures the minimum authentication interval.

The **no** form of this command returns the value to the default.

Default min-auth-interval 15

Parameters *min-auth-interval* — Specifies the minimum authentication interval, in minutes.

Values 1 to 6000

authentication-policy

Syntax **authentication-policy** *name*
no authentication-policy

Context config>service>vpls>sap

Description This command defines which subscriber authentication policy must be applied when a DHCP message is received on the interface. The authentication policies must already be defined. The policy is only applied when DHCP snooping is enabled on the SAP.

arp-reply-agent

Syntax **arp-reply-agent** [**sub-ident**]

no arp-reply-agent

Context config>service>vpls>sap

Description This command enables a special ARP response mechanism in the system for ARP requests destined to static or dynamic hosts associated with the SAP. The system responds to each ARP request using the hosts MAC address as the both the source MAC address in the Ethernet header and the target hardware address in the ARP header.

ARP replies and requests received on a SAP with **arp-reply-agent** enabled are evaluated by the system against the anti-spoof filter entries associated with the ingress SAP (if the SAP has anti-spoof filtering enabled). ARPs from unknown hosts on the SAP are discarded when anti-spoof filtering is enabled.

The ARP reply agent only responds if the ARP request enters an interface (SAP, spoke-SDP or mesh-SDP) associated with the VPLS instance of the SAP.

A received ARP request that is not in the ARP reply agent table is flooded to all forwarding interfaces of the VPLS capable of broadcast except the ingress interface while honoring split-horizon constraints.

Static hosts can be defined on the SAP using the **host** command. Dynamic hosts are enabled on the system by enabling the **lease-populate** command in the SAP's **dhcp** context. In the event that both a static host and a dynamic host share the same IP and MAC address, the VPLS ARP reply agent retains the host information until both the static and dynamic information are removed. In the event that both a static and dynamic host share the same IP address, but different MAC addresses, the VPLS ARP reply agent is populated with the static host information.

The **arp-reply-agent** command fails if an existing static host on the SAP does not have both MAC and IP addresses specified. Once the ARP reply agent is enabled, creating a static host on the SAP without both an IP address and MAC address fails.

The ARP-reply-agent may only be enabled on SAPs supporting Ethernet encapsulation.

The **no** form of this command disables ARP-reply-agent functions for static and dynamic hosts on the SAP.

Parameters **sub-ident** — Configures the arp-reply-agent to discard ARP requests received on the SAP that are targeted for a known host on the same SAP with the same subscriber identification.

Hosts are identified by their subscriber information. For DHCP subscriber hosts, the subscriber hosts, the subscriber information is configured using the optional subscriber parameter string.

When arp-reply-agent is enabled with **sub-ident**:

- If the subscriber information for the destination host exactly matches the subscriber information for the originating host and the destination host is known on the same SAP as the source, the ARP request is silently discarded.

- If the subscriber information for the destination host or originating host is unknown or undefined, the source and destination hosts are not considered to be the same subscriber. The ARP request is forwarded outside the SAP's Split Horizon Group.
- When **sub-ident** is not configured, the arp-reply-agent does not attempt to identify the subscriber information for the destination or originating host and does not discard an ARP request based on subscriber information.

bpdu-translation

Syntax **bpdu-translation** {**auto** | **auto-rw** | **pvst** | **pvst-rw** | **stp**}
no bpdu-translation

Context config>service>vpls>sap
config>service>vpls>spoke-sdp

Description This command enables the translation of BPDUs to a given format, meaning that all BPDUs transmitted on a given SAP or spoke SDP have a specified format.

The **no** form of this command reverts to the default setting.

Parameters **auto** — Specifies that appropriate format is detected automatically, based on type of BPDUs received on such port.

pvst-rw — Specifies the BPDU format as PVST and the rewrite of the VLAN-ID occurs when:

- The BPDU sent on the dot1Q SAP egress contains the VLAN ID of the SAP in the BDPU-PVID TLV
- The BPDU sent on the default QinQ SAP egress contains the outer VLAN ID of the SAP in the BDPU-PVID TLV
- The BPDU sent on the QinQ SAP egress contains the inner VLAN ID of the SAP in the BDPU-PVID TLV

pvst — Specifies the BPDU format as PVST.

Note: The correct VLAN tag is included in the payload (depending on encapsulation value of outgoing SAP).

auto-rw — Specifies that the appropriate format is detected automatically and the rewrite of the VLAN-ID occurs when:

- The BPDU sent on the dot1Q SAP egress contains the VLAN ID of the SAP in the BDPU-PVID TLV
- The BPDU sent on the default QinQ SAP egress contains the outer VLAN ID of the SAP in the BDPU-PVID TLV

- The BPDU sent on the QinQ SAP egress contains the inner VLAN ID of the SAP in the BDPU-PVID TLV

stp — Specifies the BPDU-format as STP.

mfib-table-high-wmark

Syntax [**no**] **mfib-table-high-wmark** *high-water-mark*

Context config>service>vpls

Description This command specifies the multicast FIB high watermark. When the percentage filling level of the multicast FIB exceeds the configured value, a trap is generated and/or a log entry is added.

The **no** form of this command reverts to the default.

Default mfib-table-high-wmark 95

Parameters *high-water-mark* — Specifies the multicast FIB high watermark as a percentage.

Values 1 to 100

mfib-table-low-wmark

Syntax [**no**] **mfib-table-low-wmark** *low-water-mark*

Context config>service>vpls

Description This command specifies the multicast FIB low watermark. When the percentage filling level of the Multicast FIB drops below the configured value, the corresponding trap is cleared and/or a log entry is added.

The **no** form of this command reverts to the default.

Default mfib-table-low-wmark 90

Parameters *low-water-mark* — Specifies the multicast FIB low watermark as a percentage.

Values 1 to 100

mfib-table-size

Syntax **mfib-table-size** *size*
no mfib-table-size

Context config>service>vpls

Description This command specifies the maximum number of (s,g) entries in the multicast forwarding database (MFIB) for this VPLS instance.

The *mfib-table-size* parameter specifies the maximum number of multicast database entries for both learned and static multicast addresses for the VPLS instance.

When a table-size limit is set on the mfib of a service which is lower than the current number of dynamic entries present in the mfib then the number of entries remains above the limit.

The **no** form of this command removes the configured maximum MFIB table size.

Parameters *size* — Specifies the maximum number of (s,g) entries allowed in the Multicast FIB.

Values 1 to 16383

send-flush-on-failure

Syntax [**no**] **send-flush-on-failure**

Context config>service>vpls

Description This command enables sending out “flush-all-from-ME” messages to all LDP peers included in affected VPLS, in the event of physical port failures or “oper-down” events of individual SAPs. This feature provides an LDP-based mechanism for recovering a physical link failure in a dual-homed connection to a VPLS service. This method provides an alternative to RSTP solutions where dual homing redundancy and recovery, with link failure, is resolved by RSTP running between a PE router and CE devices.

This feature cannot be enabled on management VPLS.

The **no** form of this command reverts to the default.

2.11.2.2.2 Service DHCP and Anti-Spoof Filtering Commands

anti-spoof

Syntax **anti-spoof** *type*
no anti-spoof

Context config>service>vpls>sap
config>service>ies>sap

Description This command enables anti-spoof filtering and optionally changes the anti-spoof matching type for the SAP.

The type of anti-spoof filtering defines what information in the incoming packet is used to generate the criteria to lookup an entry in the anti-spoof filter table. The type parameter (**ip**, **mac**, **ip-mac**) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled.

The **no** form of this command disables anti-spoof filtering on the SAP.

Parameters *type* — Specifies the anti-spoof filtering type for this SAP.

Values ip — Specifies SAP anti-spoof filtering to use only the source IP address in its lookup. If a static host exists on the SAP without an IP address specified, the **anti-spoof ip** command fails.

mac — Specifies SAP anti-spoof filtering to use only the source MAC address in its lookup. If a static host exists on the SAP without a specified MAC address, the **anti-spoof mac** command fails.

ip-mac — Specifies SAP anti-spoof filtering to use both the source IP address and the source MAC address in its lookup.

2.11.2.2.3 Interface IPv6 Commands

address

Syntax **address** *ipv6-address/prefix-length* [**eui-64**]
no address *ipv6-address/prefix-length*

Context config>service>ies>if>ipv6

Description This command assigns an IPv6 address to the IES interface.

The **no** form of this command reverts to the default.

Parameters *ipv6-address/prefix-length* — Specifies the IPv6 address on the interface.

Values

ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x [0 to FFFF]H d [0 to 255]D
prefix-length	1 to 128

eui-64 — When the **eui-64** keyword is specified, a complete IPv6 address from the supplied prefix and 64-bit interface identifier is formed. The 64-bit interface identifier is derived from MAC address on Ethernet interfaces. For interfaces without a MAC address, for example ATM interfaces, the Base MAC address of the chassis is used.

2.11.2.3 Show Commands

Note: The command outputs in this section are examples only; actual displays may differ depending on supported functionality and user configuration

active-subscribers

Syntax **active-subscribers detail**
active-subscribers mirror
active-subscribers summary

Context show>service

Description This command displays active subscriber information.

Parameters **summary** — Displays active subscriber information in a brief format.

detail — Displays detailed output.

Output The following output is an example of active service subscribers information.

Sample Output

```
*A:Dut-C# show service active-subscribers
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber hpolSub1 (hpolSubProf2)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:lag-1:2000.1 - sla:hpolSlaProf1
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
239.1.4.194
                00:01:00:00:03:c1    PPP 1          IPCP         1000       Y
-------------------------------------------------------------------------------
(2) SLA Profile Instance sap:lag-1:2000.1 - sla:hpolSlaProf2
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
239.1.4.35
                00:01:00:00:03:22    N/A            ARP-Host     1000       Y
239.1.4.195
                00:01:00:00:03:c2    PPP 1          IPCP         1000       Y
-------------------------------------------------------------------------------
Subscriber hpolSub16 (hpolSubProf1)
-------------------------------------------------------------------------------
```

```
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[lag-1:2000.2] - sla:hpolSlaProf1
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
239.1.4.224
                00:01:00:00:03:df    PPP 1          IPCP         1000       Y
-------------------------------------------------------------------------------
Subscriber hpolSub2 (hpolSubProf1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:lag-1:2000.1 - sla:hpolSlaProf1
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
239.1.4.196
                00:01:00:00:03:c3    PPP 1          IPCP         1000       Y
-------------------------------------------------------------------------------
Number of active subscribers : 3
===============================================================================
*A:Dut-C#
```

Table 12 describes active subscriber output fields.

Table 12 **Active Subscriber Field Descriptions**

Field	Description
Active Subscribers	The active subscriber's name
IP Address	The IP address of the active subscriber
MAC Address	The MAC address of the active subscriber
Session	The session type
Origin	The IPCP parameter
Svc	The service ID
Fwd	Option to forward
Number of active subscribers	The total number of active subscribers

credit-control

Syntax **credit-control** [**subscriber** *sub-ident-string*]
credit-control out-of-credit [**action** *action*] [**summary**]
credit-control extended-failure-handling [*state*] [**summary**]

Context show>service>active-subscribers

Description This command displays active subscriber credit control information. Without additional filters, the output lists information for all active subscribers.

The output of this command can be filtered:

- per subscriber by specifying the optional **subscriber** *sub-ident-string* parameter
- per all subscribers that are out of credit by specifying the **out-of-credit** parameter with the optional **action** *action* parameter
- per subscribers that have a Diameter Gy session with Extended Failure Handling (EFH) enabled by specifying the **extended-failure-handling** parameter with the optional *state* parameter.

Parameters *sub-ident-string* — Displays credit control information for the specified subscriber, up to 32 characters.

out-of-credit [**action** *action*] — Displays credit control information for subscribers that are out of credit and, optionally, have the specified out of credit action enabled.

Values continue, block-category, change-service-level

extended-failure-handling [*state*] — Displays credit control information for subscribers with a Diameter Gy session with EFH enabled. The optional *state* parameter lists subscribers for which the Diameter Gy session EFH is active, inactive, or both (all).

Values active, inactive, all

summary — Lists the number of matching entries for the specified filter.

Output The following output is an example of active subscribers credit control information.

Output Sample

```
A:BNG-1# show service active-subscribers credit-control
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber ipoe-msap-002 (sub-profile-1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[1/1/4:1201.2] - sla:sla-profile-3
-------------------------------------------------------------------------------
Credit Control Policy: cc-policy-1
Category Map         : cat-map-3
Diameter Session Gy  : bng.domain.com;1464610029;840
CC Failure Handling  : continue
Extended Failure Handling (EFH)
  State              : inactive
  Total Active time  : 0d 00:10:02
  Total Active Count : 1
Number of categories
static               : 3
gx-session           : 0
gx-pcc               : 0
Category Name        : cat-1-time
Ingress Queues       : 1
```

```
Egress Queues        : 1
Ingress Policers     :
Egress Policers      :
Credit Volume Used   : 0                    Credit Time Used     : 39
Credit Volume Avail. : 0                    Credit Time Avail.   : 861
Credit Volume Thres. : 0                    Credit Time Thres.   : 0
Credit Expired       : False                Credit Negotiating   : False
Out Of Credit Action : None                 Quota Holding Time   : 0
Validity Time Used   : 0                    Validity Time Avail. : 0
EFH Unreported Credit
  Total Volume       : 0                      Total Time         : 1195
HTTP Rdr URL Override: (Not Specified)
Category Name        : cat-1-volume
Ingress Queues       : 1
Egress Queues        : 1
Ingress Policers     :
Egress Policers      :
Credit Volume Used   : 31800                Credit Time Used     : 0
Credit Volume Avail. : 4968200              Credit Time Avail.   : 0
Credit Volume Thres. : 0                    Credit Time Thres.   : 0
Credit Expired       : False                Credit Negotiating   : False
Out Of Credit Action : None                 Quota Holding Time   : 0
Validity Time Used   : 32                   Validity Time Avail. : 568
EFH Unreported Credit
  Total Volume       : 1110400                Total Time         : 0
HTTP Rdr URL Override: (Not Specified)
Category Name        : cat-2-volume
Ingress Queues       : 4
Egress Queues        : 4
Ingress Policers     :
Egress Policers      :
Credit Volume Used   : 98400                Credit Time Used     : 0
Credit Volume Avail. : 1901600              Credit Time Avail.   : 0
Credit Volume Thres. : 0                    Credit Time Thres.   : 0
Credit Expired       : False                Credit Negotiating   : False
Out Of Credit Action : None                 Quota Holding Time   : 0
Validity Time Used   : 33                   Validity Time Avail. : 567
EFH Unreported Credit
  Total Volume       : 3330600                Total Time         : 0
HTTP Rdr URL Override: (Not Specified)
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
10.1.1.101
                00:51:00:00:00:02    IPoE           DHCP         1000       Y
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Number of active subscribers : 1
===============================================================================
```

Table 13 describes credit control output fields.

(1) These fields are only shown when EFH is enabled in the Diameter application policy. EFH is enabled using the **configure subscriber-mgmt diameter-application-policy gy extended-failure-handling no shutdown** command.

(2) These fields are only shown when the EFH state is active.

Table 13 **Credit Control Field Descriptions**

Label	Description
Credit Control Policy	The credit control policy name in use for this subscriber session
Category Map	The active category map in use for this subscriber session
Diameter Session Gy	The Diameter Gy session ID
CC Failure Handling	The current value for Credit Control Failure Handling (CCFH) as configured in the Diameter Gy application policy or as received in an Answer message from the Online Charging Server (OCS) in the CCFH AVP Values: terminate, continue, or retry-and-terminate
Extended Failure Handling (EFH) - State (1)	The EFH state: active — EFH is enabled and currently active inactive — EFH is enabled and currently inactive
EFH - Attempts (1) (2)	The number of times interim credit is assigned to all rating groups followed by an attempt to establish a new Diameter Gy session with the Online Charging Server (OCS), or credit control server
EFH - Maximum Attempts(1) (2)	The configured maximum attempts value in the Diameter application policy When an attempt to establish a new Diameter Gy session with the OCS continues to fail after the maximum attempts value is reached, then the user session is terminated (the subscriber hosts are deleted from the system).
EFH - Active Time (1) (2)	The time since the EFH state became active for this subscriber session
EFH - Total Active Time(1)	The accumulated time of all occurrences that EFH was active during the lifetime of this subscriber session
EFH - Total Active Count (1)	The number of times that EFH was active during the lifetime of this subscriber session
Number of categories - static	The number of static categories in use for this subscriber session.
Number of categories - gx-session	The number of categories reserved for this subscriber session for Gx session level Usage Monitoring.

Table 13 **Credit Control Field Descriptions (Continued)**

Label	Description
Number of categories - gx-pcc	The number of dynamic categories in use by this subscriber session for GX PCC rule-based Usage Monitoring.
Category - Category Name	The name of the category as configured in the category map
Category - Ingress Queues	The ingress queues for which credit control is applied in this category as configured in the category map
Category - Egress Queues	The egress queues for which credit control is applied in this category as configured in the category map
Category - Ingress Policers	The ingress policers for which credit control is applied in this category as configured in the category map
Category - Egress Policers	The egress policers for which credit control is applied in this category as configured in the category map
Category - Credit Volume Used	The used total octets counter (ingress and egress combined) since the last usage reporting of this rating group
Category - Credit Volume Avail.	The remaining total granted octets available before a new credit negotiation is triggered for this rating group (category) or the out of credit action is triggered
Category - Credit Volume Thres.	The volume quota threshold in octets as received from the OCS. When the Credit Volume Available octets fall below the Credit Volume Threshold, a new credit negotiation is triggered for this rating group (category). The Credit Volume Threshold value is set to Expired and Credit Negotiating is set to True.
Category - Credit Time Used	The number of seconds passed since the last reporting of this rating group
Category - Credit Time Avail.	The remaining number of seconds before a new credit negotiation is triggered for this rating group (category) or the out of credit action is triggered

Table 13 **Credit Control Field Descriptions (Continued)**

Label	Description
Category - Credit Time Thres.	The time quota threshold in seconds as received from the OCS When the Credit Time Available falls below the Credit Time Threshold, a new credit negotiation is triggered for this rating group (category). The Credit Time Threshold value is set to Expired and Credit Negotiating is set to true.
Category - Credit Expired	The credit expired value: False—Negotiated credit is still available True—All granted credit is exhausted; either new credit is being negotiated with the OCS or the out of credit action is activated.
Category - Credit Negotiating	The credit negotiating value: False—There is no credit being requested to the OCS True—New credit is being requested to the OCS
Category - Out Of Credit Action	The out of credit action value: None—Out-of-credit action is not active When the out-of-credit action is active, the value is set to one of: ChangeServiceLevel, BlockCategory, Continue, or DisconnectHost
Category - Quota Holding Time	The Idle timeout associated with the granted quota. When no traffic associated with the quota is observed for the time specified, a credit negotiation is triggered with the OCS.
Category - Validity Time Used	The number of seconds passed since the Validity Time was installed
Category - Validity Time Avail.	The remaining number of seconds of the Validity Time before a new usage reporting is triggered for this rating group (category)
Category - EFH Unreported Credit - Current Volume (1) (2)	The unreported volume credit for the current occurrence of EFH in an active state. This counter includes the unreported used volume credit for the initial Diameter Gy session that caused the active EFH state and the unreported volume interim credit for previous attempts. Used interim credit for the current attempt is shown in the Credit Volume Used counter.

Table 13 **Credit Control Field Descriptions (Continued)**

Label	Description
Category - EFH Unreported Credit - Total Volume (1)	The accumulated total unreported volume credit for the previous occurrences of EFH in an active state. The total counter is updated when the EFH state toggles from active to inactive. When interim credit reporting is enabled, the counters are reset to zero when the actual usage reporting for this rating group (category) occurs. When interim credit reporting is disabled, the counters accumulate the total unreported volume credit during the lifetime of the subscriber session.
Category - EFH Unreported Credit - Current Time (1) (2)	The unreported time credit for the current occurrence of EFH in an active state. This counter includes the unreported used time credit for the initial Diameter Gy session that caused the active EFH state and the unreported time interim credit for previous attempts. Used interim credit for the current attempt is shown in the Credit Time Used counter.
Category - EFH Unreported Credit - Total Time (1)	The accumulated total unreported time credit for the previous occurrences of EFH in an active state. The total counter is updated when the EFH state toggles from active to inactive. When interim credit reporting is enabled, the counters are reset to zero when the actual usage reporting for this rating group (category) occurs. When interim credit reporting is disabled, the counters accumulate the total unreported time credit during the lifetime of the subscriber session.
Category - HTTP Rdr URL Override	The URL returned by a Diameter OCS in a Credit Control Answer (CCA) message Final-Unit-Indication/ Redirect-Server/Redirect-Server-Address AVP. This URL is used for Diameter Gy http-redirect out-of-credit actions when **allow-override** is configured.

filter

Syntax **filte**r [**subscriber** *sub-ident-string*] [**origin** *origin*]

Context show>service>active-subscribers

Description This command displays active subscriber filter information.

Parameters *sub-ident-string* — The subscriber ID of the active subscriber, up to 32 characters

origin — The origin of the active subscriber

hierarchy

Syntax **hierarchy** [**subscriber** *sub-ident-string*]

Context show>service>active-subscribers

Description This command displays active subscriber hierarchy information. To display an IPoE/PPP session, the group of hosts within the session is visually indented. Additional information related to the session is also shown. For PPPoE, the circuit ID and remote ID if used is shown. For IPoE, the key used to group the session is shown (for example, the circuit ID). The command also display PD host which are modeled as a managed route. The PD managed route is directly underneath and points to the host that it is using as the next hop. The PD managed route forwarding status is also shown, where (N) indicates that the route is not forwarding.

Parameters *sub-ident-string* — The subscriber ID of the active subscriber, up to 32 characters

Output The following output displays active subscriber hierarchy information.

Sample Output

```
*A:PE-5# show service active-subscribers hierarchy
===============================================================================
Active Subscribers Hierarchy
===============================================================================
-- ipoe-msap-001 (sub-profile-1)
   |
   +-- sap:[1/1/4:1201.1] - sla:sla-profile-1
       |
       +-- IPOE-session - mac:00:51:00:00:00:01 - svc:1000
           |
           |-- 10.1.1.29 - DHCP
           |
           +-- 2001:db8:1:100::1/128 - DHCP6
               |
               +-- 2001:db8:d001::/56 - DHCP6-PD-MR
-------------------------------------------------------------------------------
Number of active subscribers : 1
Flags: (N) = the host or the managed route is in non-forwarding state
===============================================================================

*A:PE-5# show service active-subscribers hierarchy radius-acct
===============================================================================
Active Subscribers Hierarchy (accounting information)
===============================================================================
-- ipoe-msap-001 (sub-profile-1)
   |   Acct-policy       : acct-policy-4
   |       Mode          : Queue Instance (Interim), Session (Interim)
   |       Cfg. Interval : 0180d00h
   |
   +-- sap:[1/1/4:1201.1] - sla:sla-profile-1
       |   |
       |   +-- Queue Instance Acct - acct-policy-4
       |       session-id          : 144DFF0000000358A70696
       |       multi-session-id    : 144DFF0000000358A70696
       |       interim-interval    : 0180d00h
```

```
        |
        +-- IPOE-Session - mac:00:51:00:00:00:01
            |   |
            |   +-- Session Acct - acct-policy-4
            |       session-id      : 144DFF0000000158A70696
            |       multi-session-id: 144DFF0000000358A70696
            |       interim-interval: 0180d00h
            |
            |-- 10.1.1.29
            |   |
            |   +-- no radius acct
            |       session-id      : 144DFF0000000258A70696
            |
            +-- 2001:db8:1:100::1/128
                |
                +-- no radius acct
                |   session-id      : 144DFF0000000458A70B2B
                |
                +-- 2001:db8:d001::/56
                    session-id      : N/A
                    multi-session-id: N/A
-------------------------------------------------------------------------------
Number of active subscribers : 1
===============================================================================
```

host-tracking

Syntax **host-tracking** [**subscriber** *sub-ident-string*]
host-tracking [**subscriber** *sub-ident-string*] **detail**
host-tracking [**subscriber** *sub-ident-string*] **summary**
host-tracking [**subscriber** *sub-ident-string*] **statistics**

Context show>service>active-subscribers

Description This command displays active subscriber host tracking information.

Parameters *sub-ident-string* — Displays subscriber ID of the active subscriber, up to 32 characters.

detail — Displays detailed host tracking information for the specified subscriber ID.

summary — Displays summarized host tracking information for the specified subscriber ID.

statistics — Displays statistical host tracking information for the specified subscriber ID.

groups

Syntax **groups** [**group** *group-ip-address*]
groups [**group** *group-ip-address*] **detail**
groups [**group** *group-ip-address*] **summary**

Context show>service>active-subscribers

Description This command displays active subscriber host tracking groups information.

Parameters *group-ip-address* — Displays information about the group address of the active subscriber group

detail — Displays detailed information about the group address of the active subscriber group

summary — Displays summarized group information for active subscribers.

igmp

Syntax **igmp** [**subscriber** *sub-ident-string*][**detail**]

Context show>service>active-subscribers

Description This command displays active subscriber IGMP information.

Parameters *sub-ident-string* — Displays the subscriber ID of the active subscriber, up to 32 characters.

detail — Displays detailed IGMP information for the specified subscriber ID.

Output The following output is an example of IGMP information for active subscribers.

Sample Output

```
*B:Dut-C# show service active-subscribers igmp
===============================================================================
Active Subscribers
===============================================================================
Subscriber                      IGMP-Policy
  HostAddr                      GrpItf                          NumGroups
-------------------------------------------------------------------------------
sub_1                           pol1
  10.112.1.1                    gi_1_1                          1
  10.112.1.2                    gi_1_1                          2
sub_2                           pol1
  10.112.1.3                    gi_1_2                          0
-------------------------------------------------------------------------------
Number of Subscribers : 2
===============================================================================
*B:Dut-C#

*B:Dut-C# show service active-subscribers igmp detail
===============================================================================
Active Subscribers Detail
===============================================================================
Subscriber                      IGMP-Policy
  HostAddr                      GrpItf                          NumGroups
    GrpAddr                     Type            Up-Time         Mode
      SrcAddr                   Type                            Blk/Fwd
-------------------------------------------------------------------------------
sub_1                           pol1
  10.112.1.1                    gi_1_1                          1
```

```
    239.0.0.1                   Dynamic         0d 00:00:53      Include
      10.11.0.1                 Dynamic                          Fwd
      10.11.0.2                 Dynamic                          Fwd
  10.112.1.2                    gi_1_1                           2
    239.0.0.1                   Dynamic         0d 00:00:44      Exclude
      10.11.0.1                 Dynamic                          Blk
    239.0.0.2                   Dynamic         0d 00:00:44      Exclude
      10.11.0.1                 Dynamic                          Blk
sub_2                           pol1
  10.112.1.3                    gi_1_2                           0
-------------------------------------------------------------------------------
Number of Subscribers : 2
===============================================================================
*B:Dut-C#

*B:Dut-C# show service active-subscribers igmp subscriber "sub_1" detail
===============================================================================
Active Subscribers Detail
===============================================================================
Subscriber                      IGMP-Policy
  HostAddr                      GrpItf                           NumGroups
    GrpAddr                     Type            Up-Time          Mode
      SrcAddr                   Type                             Blk/Fwd
-------------------------------------------------------------------------------
sub_1                           pol1
  10.112.1.1                    gi_1_1                           1
    239.0.0.1                   Dynamic         0d 00:01:04      Include
      10.11.0.1                 Dynamic                          Fwd
     10.11.0.2                  Dynamic                          Fwd
  10.112.1.2                    gi_1_1                           2
    239.0.0.1                   Dynamic         0d 00:00:55      Exclude
      10.11.0.1                 Dynamic                          Blk
    239.0.0.2                   Dynamic         0d 00:00:55      Exclude
      10.11.0.1                 Dynamic                          Blk
-------------------------------------------------------------------------------
Number of Subscribers : 1
===============================================================================
B:Dut-C#
```

subscriber

Syntax **subscriber** *sub-ident-string*

subscriber *sub-ident-string* **detail**

subscriber *sub-ident-string* **mirror**

subscriber *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]

subscriber *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*] **detail**

subscriber *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*] **mirror**

Context show>service>active-subscribers

Description This command displays active subscriber information for a subscriber.

Parameters *sub-ident-string* — The subscriber ID of the active subscriber, up to 32 characters

sap-id — Displays the physical port identifier portion of the SAP definition.

sla-profile-name — Displays the existing SLA profile name, up to 32 characters.

spi-sharing-type — Displays subscriber information for sessions with the specified SPI sharing type.

Values per-group, per-session-ipoe, per-session-ppp

spi-sharing-id — Displays subscriber information for sessions with the specified SPI sharing type and sharing identifier. The SPI group ID is:

- a group id when sharing type is per-group
- an internal PPP or IPoE sharing id when the sharing type is per-PPP or per-IPoE session

Values 0 to 4294967295

detail — Displays detailed output.

mirror — Displays mirrored output.

Output The following output is an example of subscriber information for active subscribers.

Sample Output

```
*A:Dut-C# show service active-subscribers subscriber "hpolSub1"
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber hpolSub1 (hpolSubProf2)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:lag-1:2000.1 - sla:hpolSlaProf1
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
239.1.4.194
                00:01:00:00:03:c1    PPP 1          IPCP         1000       Y
-------------------------------------------------------------------------------
(2) SLA Profile Instance sap:lag-1:2000.1 - sla:hpolSlaProf2
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
239.1.4.35
                00:01:00:00:03:22    N/A            ARP-Host     1000       Y
239.1.4.195
                00:01:00:00:03:c2    PPP 1          IPCP         1000       Y
===============================================================================
*A:Dut-C#
*A:Dut-C# show service active-subscribers subscriber "hpolSub1"
sap lag-1:2000.1 sla-profile "hpolSlaProf2"
===============================================================================
Active Subscribers
```

```
===============================================================================
-------------------------------------------------------------------------------
Subscriber hpolSub1 (hpolSubProf2)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:lag-1:2000.1 - sla:hpolSlaProf2
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
239.1.4.35
                00:01:00:00:03:22    N/A            ARP-Host     1000       Y
239.1.4.195
                00:01:00:00:03:c2    PPP 1          IPCP         1000       Y
===============================================================================
```

id

Syntax **id** *service-id*

Context show>service

Description This command enables the context to display information for a particular service-id.

Parameters *service-id* — Displays information about the specified service identification that identifies the service in the service domain.

all — Displays detailed information about the service

arp — Displays ARP entries for the service

arp-host — Displays ARP host related information

base — Displays basic service information

fdb — Displays FDB entries

host — Displays static hosts configured on the specified service

labels — Displays labels being used by this service

sap — Displays SAPs associated to the service

sdp — Displays SDPs associated with the service

split-horizon-group — Display split horizon group information

stp — Displays STP information

arp

Syntax **arp** [*ip-address*] | [**mac** *ieee-address*] | [**sap** *sap-id*] | [**interface** *ip-int-name*]

Context show>service>id

Description This command displays the ARP table for the IES instance.

Parameters *ip-address* — Displays only ARP entries in the ARP table with the specified IP address

Default All IP addresses

ieee-address — Displays only ARP entries in the ARP table with the specified 48-bit MAC address. The MAC address can be expressed in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.

Default All MAC addresses.

sap-id — Displays SAP information for the specified SAP ID

ip-int-name — Displays the IP interface name for which to display matching ARPs

Output The following output is an example of ARP service ID information.

Sample Output

```
A:ALA-12# show service id 2 arp
===============================================================================
ARP Table
===============================================================================
IP Address      MAC Address       Type    Age       Interface         Port
-------------------------------------------------------------------------------
10.11.1.1       00:03:fa:00:08:22 Other   00:00:00  ies-100-10.11.1  1/1/11:0
===============================================================================
A:ALA-12#
```

Table 14 describes show service-id ARP output fields.

Table 14 **Show Service-ID ARP Field Descriptions**

Label	Description
Service ID	The value identifying the service
MAC	The specified MAC address
Source-Identifier	The location the MAC is defined
Type	Static — FDB entries created by management. Learned — Dynamic entries created by the learning process. OAM — The entries created by the OAM process.
Age	The time lapsed since the service was enabled
Interface	The interface applied to the service
Port	The port where the SAP is applied

base

Syntax **base**

Context show>service>id

Description This command displays basic information about the service ID including service type, description, SAPs and SDPs.

Output The following output is an example of base service ID information.

Sample Output

```
*A:PE-6# show service id 1 base

===============================================================================
Service Basic Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPLS
Name              : 1
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 05/08/2018 09:40:32
Last Mgmt Change  : 05/08/2018 09:40:24
Etree Mode        : Disabled
Admin State       : Up                  Oper State        : Up
MTU               : 1514
SAP Count         : 1                   SDP Bind Count    : 1
Snd Flush on Fail : Disabled            Host Conn Verify  : Disabled
SHCV pol IPv4     : None
Propagate MacFlush: Disabled            Per Svc Hashing   : Disabled
Allow IP Intf Bind: Disabled
Fwd-IPv4-Mcast-To*: Disabled            Fwd-IPv6-Mcast-To*: Disabled
Mcast IPv6 scope  : mac-based
Def. Gateway IP   : None
Def. Gateway MAC  : None
Temp Flood Time   : Disabled            Temp Flood        : Inactive
Temp Flood Chg Cnt: 0
SPI load-balance  : Disabled
TEID load-balance : Disabled
Src Tep IP        : N/A
Vxlan ECMP        : Disabled
VSD Domain        : <none>

-------------------------------------------------------------------------------
Service Access & Destination Points
-------------------------------------------------------------------------------
Identifier                              Type         AdmMTU  OprMTU  Adm   Opr
-------------------------------------------------------------------------------
sap:1/1/c1/1:1                          q-tag        9000    9000    Up    Up
sdp:65:1 S(192.0.2.5)                   Spok         0       8974    Up    Down
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

Table 15 describes show service-id base output fields:

Table 15 **Show Service-ID Base Field Descriptions**

Label	Description
Service Id	The service identifier
Vpn Id	The VPN ID assigned to the service.
Service Type	The type of service
Description	Generic information about the service
Customer Id	The customer identifier
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer
Adm	The administrative state of the service
Oper	The operating state of the service
Mtu	The largest frame size (in octets) that the service can handle.
SAP Count	The number of SAPs defined on the service
SDP Bind Count	The number of SDPs bound to the service
Identifier	The service access (SAP) and destination (SDP) points
Type	The signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on the SDP.
AdmMTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end ESR, without requiring the packet to be fragmented.
OprMTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end ESR, without requiring the packet to be fragmented
Opr	The operational state of the SDP

authentication

Syntax **authentication**

Context show>service>id

Description This command enables the context to show session authentication information.

statistics

Syntax **statistics** [**policy** *name*] [**sap** *sap-id*]

Context show>service>id>authentication

Description This command displays session authentication statistics for this service.

Parameters **policy** *name* — Specifies an existing authentication policy name.

sap-id — Specifies the physical port identifier portion of the SAP definition.

Output The following output displays an example of service authentication statistics information

Sample Output

```
*A:ALA-48# show service id 700 authentication statistics
===============================================================
Authentication Statistics for service 700
===============================================================
Client Packets Authenticate Fail  : 0
Client Packets Authenticate Ok    : 0
===============================================================
*A:ALA-48#
```

egress-label

Syntax **egress-label** *egress-label1* [*egress-label2*]

Context show>service

Description This command displays services using the range of egress labels.

If only the mandatory *egress-label1* parameter is specified, only services using the specified label are displayed.

If both *egress-label1* and *egress-label2* parameters are specified, the services using the range of labels X where *egress-label1* <= X <= *egress-label2* are displayed.

Use the **show router ldp bindings** command to display dynamic labels.

Parameters *egress-label1* — Specifies the starting egress label value for which to display services using the label range. If only *egress-label1* is specified, services only using *egress-label1* are displayed.

Values 0, 2049 to 131071

egress-label2 — Specifies the ending egress label value for which to display services using the label range.

Values 2049 to 131071

Default The *egress-label1* value

Output The following output is an example of service egress label information.

Sample Output

```
A:ALA-12# show service egress-label 0 10000
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Id             Type I.Lbl        E.Lbl
-------------------------------------------------------------------------------
1          10:1               Mesh 0            0
1          20:1               Mesh 0            0
1          30:1               Mesh 0            0
1          100:1              Mesh 0            0
...
1          107:1              Mesh 0            0
1          108:1              Mesh 0            0
1          300:1              Mesh 0            0
1          301:1              Mesh 0            0
1          302:1              Mesh 0            0
1          400:1              Mesh 0            0
1          500:2              Spok 131070       2001
1          501:1              Mesh 131069       2000
100        300:100            Spok 0            0
200        301:200            Spok 0            0
300        302:300            Spok 0            0
400        400:400            Spok 0            0
-------------------------------------------------------------------------------
Number of Bindings Found : 23
===============================================================================
A:ALA-12#
```

The following table describes show service egress label output fields.

Table 16 **Service Egress Label Output Field Descriptions**

Label	Description
Svc Id	The ID that identifies a service
Sdp Id	The ID that identifies an SDP
Type	Indicates whether the SDP binding is a spoke or a mesh
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP
E. Lbl	The VC label used by this device to send packets to the far-end device in this service by the SDP
Number of bindings found	The total number of SDP bindings that exist within the specified egress label range

fdb-info

Syntax **fdb-info**

Context show>service

Description This command displays global FDB usage information.

Output The following output is an example of service FDB information.

Sample Output

```
A:ALA-12# show service fdb-info
===============================================================================
Forwarding Database (FDB) Information
===============================================================================
Service Id        : 700                 Mac Move          : Disabled
Mac Move Rate     : 10                  Mac Move Timeout  : 10
Table Size        : 250                 Total Count       : 0
Learned Count     : 0                   Static Count      : 0
Remote Age        : 900                 Local Age         : 300
High WaterMark    : 95%                 Low Watermark     : 90%
Mac Aging         : Enabl               Relearn Only      : False

Service Id        : 725                 Mac Move          : Disabled
Mac Move Rate     : 10                  Mac Move Timeout  : 10
Table Size        : 250                 Total Count       : 0
Learned Count     : 0                   Static Count      : 0
Remote Age        : 900                 Local Age         : 300
High WaterMark    : 95%                 Low Watermark     : 90%
Mac Learning      : Enabl               Discard Unknown   : Dsabl
Mac Aging         : Enabl               Relearn Only      : False

Service Id        : 740                 Mac Move          : Disabled
Mac Move Rate     : 10                  Mac Move Timeout  : 10
Table Size        : 250                 Total Count       : 0
Learned Count     : 0                   Static Count      : 0
Remote Age        : 900                 Local Age         : 300
High WaterMark    : 95%                 Low Watermark     : 90%
Mac Learning      : Enabl               Discard Unknown   : Dsabl
Mac Aging         : Enabl               Relearn Only      : False
...
-------------------------------------------------------------------------------
Total Service FDBs : 7
Total FDB Configured Size : 1750
Total FDB Entries In Use  : 0
-------------------------------------------------------------------------------
===============================================================================
*A:ALA-48#
```

The following table describes show FDB-Info command output.

Table 17 **Show FDB-Info Command Output Field Descriptions**

Label	Description
Service ID	The service ID
Mac Move	The administrative state of the MAC movement feature associated with the service
Mac Move Rate	The maximum rate at which MACs can be re-learned in this service, before the SAP where the moving MAC was last seen is automatically disabled in order to protect the system against undetected loops or duplicate MACs. The rate is computed as the maximum number of re-learns allowed in a 5 second interval: for example, the default rate of 10 re-learns per second corresponds to 50 re-learns in a 5 second period.
Mac Move Timeout	The time, in seconds, to wait before a SAP that has been disabled after exceeding the maximum re-learn rate is re-enabled. A value of zero indicates that the SAP is not automatically re-enabled after being disabled. If after the SAP is re-enabled it is disabled again, the effective retry timeout is doubled in order to avoid thrashing.
Table Size	The maximum number of learned and static entries allowed in the FDB
Total Count	The current number of entries (both learned and static) in the FDB
Learned Count	The current number of learned entries in the FDB
Static Count	The current number of static entries in the FDB
Remote Age	The number of seconds used to age out FDB entries learned on an SDP. These entries correspond to MAC addresses learned on remote SAPs.
Local Age	The number of seconds used to age out FDB entries learned on local SAPs
High WaterMark	The utilization of the FDB table of this service at which a ‘table full’ alarm is raised by the agent
Low WaterMark	The utilization of the FDB table of this service at which a ‘table empty’ alarm is raised by the agent
Mac Learning	Specifies whether the MAC learning process is enabled in this service
Discard Unknown	Specifies whether frames received with an unknown destination MAC are discarded in this service
MAC Aging	Specifies whether the MAC aging process is enabled in this service

Table 17 **Show FDB-Info Command Output Field Descriptions (Continued)**

Label	Description (Continued)
MAC Pinning	Specifies whether MAC Pinning is enabled in this service
Relearn Only	When enabled, indicates that either the FDB table of this service is full or that the maximum system-wide number of MACs supported by the agent has been reached, and thus MAC learning is temporary disabled, and only MAC re-learns can take place.
Total Service FDBs	The current number of service FDBs configured on this node
Total FDB Size	The sum of configured FDBs
Total FDB Entries In Use	The total number of entries (both learned and static) in use

fdb-mac

Syntax **fdb-mac** *ieee-address*

Context show>service

Description This command displays the FDB entry for a given MAC address.

Parameters *ieee-address* — Specifies the 48-bit MAC address for which to display the FDB entry in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.

Output The following output is an example of service FDB MAC information.

Sample Output

```
A:ALA-12# show service fdb-mac 00:99:00:00:00:00
===============================================================================
Services Using Forwarding Database Mac 00:99:00:00:00:00
===============================================================================
ServId  MAC                       Source-Identifier         Type/
Age Last Change
-------------------------------------------------------------------------------
1       00:99:00:00:00:00         sap:1/2/7:0               Static
===============================================================================
A:ALA-12#
```

The following table describes the show FDB MAC command output fields:

Table 18 **Show FDB-MAC Command Output Field Descriptions**

Label	Description
Service ID	The value that identifies a specific service
MAC	The specified MAC address
Source-Identifier	The location where the MAC is defined
Type	Static — The FDB entries created by management. Learned — The dynamic entries created by the learning process. OAM — The entries created by the OAM process.

ingress-label

Syntax **ingress-label** *ingress-label1* [*ingress-label2*]

Context show>service

Description This command displays service information using the range of ingress labels.

If only the mandatory *ingress-label1* parameter is specified, only services using the specified label are displayed.

If both *ingress-label1* and *ingress-label2* parameters are specified, the services using the range of labels X where *ingress-label1* <= X <= *ingress-label2* are displayed.

Use the **show router vprn-service-id ldp bindings** command to display dynamic labels.

Parameters *ingress-label1* — Displays the starting ingress label value for which to display services using the label range. If only *ingress-label1* is specified, services only using *ingress-label1* are displayed.

Values 0, 2048 to 131071

ingress-label2 — Displays the ending ingress label value for which to display services using the label range.

Values 2048 to 131071

Default The *ingress-label1* value.

Output The following output is an example of service ingress label information.

Sample Output

```
A:ALA-12# show service ingress label 0
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Id             Type I.Lbl        E.Lbl
```

```
-------------------------------------------------------------------------------
1           10:1                Mesh 0            0
1           20:1                Mesh 0            0
1           30:1                Mesh 0            0
1           50:1                Mesh 0            0
1           100:1               Mesh 0            0
1           101:1               Mesh 0            0
1           102:1               Mesh 0            0
1           103:1               Mesh 0            0
1           104:1               Mesh 0            0
1           105:1               Mesh 0            0
1           106:1               Mesh 0            0
1           107:1               Mesh 0            0
1           108:1               Mesh 0            0
1           300:1               Mesh 0            0
1           301:1               Mesh 0            0
1           302:1               Mesh 0            0
1           400:1               Mesh 0            0
1           500:2               Spok 131070       2001
1           501:1               Mesh 131069       2000
100         300:100             Spok 0            0
200         301:200             Spok 0            0
300         302:300             Spok 0            0
400         400:400             Spok 0            0
-------------------------------------------------------------------------------
Number of Bindings Found : 23
-------------------------------------------------------------------------------
A:ALA-12#
```

The following table describes show service ingress label output fields.

Table 19 **Show Service Ingress Label Field Descriptions**

Label	Description
Svc ID	The value that identifies a specific service
SDP Id	The SDP identifier
Type	Indicates whether the SDP is a spoke or a mesh
I.Lbl	The ingress label used by the far-end device to send packets to this device in this service by the SDP
E.Lbl	The egress label used by this device to send packets to the far-end device in this service by the SDP
Number of Bindings Found	The number of SDP bindings within the label range specified

sap-using

Syntax **sap-using** [**sap** *sap-id*]

sap-using interface [*ip-address* | *ip-int-name*]
sap-using [**ingress** | **egress**] **atm-td-profile** *td-profile-id*
sap-using [**ingress** | **egress**] **filter** *filter-id*
sap-using [**ingress** | **egress**] **qos-policy** *qos-policy-id*
sap-using authentication-policy *auth-plcy-name*

Context show>service

Description This command displays SAP information.

If no optional parameters are specified, the command displays a summary of all defined SAPs.

The optional parameters restrict output to only SAPs matching the specified properties.

Parameters **ingress** — Specifies matching an ingress policy.

egress — Specifies matching an egress policy.

qos-policy-id — Displays the ingress or egress QoS Policy ID for which to display matching SAPs.

Values 1 to 65535

td-profile-id — Displays SAPs using this traffic description (applies only to the 7750 SR).

filter-id — Displays the ingress or egress filter policy ID for which to display matching SAPs.

Values 1 to 65535

auth-plcy-name — Displays the session authentication policy for which to display matching SAPs.

sap-id — Specifies the physical port identifier portion of the SAP definition.

interface — Specifies matching SAPs with the specified IP interface.

ip-addr — Displays the IP address of the interface for which to display matching SAPs.

Values 1.0.0.0 to 223.255.255.255

ip-int-name — Displays the IP interface name for which to display matching SAPs.

Output The following output is an example of service SAP information.

Sample Output

```
A:ALA-12# show service sap-using sap 1/1
===============================================================================
Service Access Points
===============================================================================
PortId          SvcId      SapMTU I.QoS I.Mac/IP E.QoS E.Mac/IP A.Pol  Adm  Opr
-------------------------------------------------------------------------------
1/1/7:0         1          1518   10    8        10    none     none   Up   Up
1/1/11:0        100        1514   1     none     1     none     none   Down Down
1/1/7:300       300        1518   10    none     10    none     1000   Up   Up
```

```
-------------------------------------------------------------------------------
Number of SAPs : 3
-------------------------------------------------------------------------------
A:ALA-12#
```

The following example applies to the 7450 ESS:

```
ALA-1# show service sap-using authentication-policy dhcp-radius-plcy
===============================================================================
Service Access Point Using Authentication Policy dhcp-radius-plcy
===============================================================================
PortId                SvcId      I.QoS I.Fltr E.QoS E.Fltr Acc.Pol Adm  Opr
-------------------------------------------------------------------------------
1/1/1                 1          1     none   1     none   none    Up   Down
-------------------------------------------------------------------------------
Saps : 1
-------------------------------------------------------------------------------
A:ALA-1#
```

The following table describes show service SAP output fields:

Table 20 **Show Service SAP Field Descriptions**

Label	Description
Port ID	The ID of the access port where the SAP is defined
Svc ID	The service identifier
SapMTU	The SAP MTU value
I.QoS	The SAP ingress QoS policy number specified on the ingress SAP
I.MAC/IP	The MAC or IP filter policy ID applied to the ingress SAP
E.QoS	The SAP egress QoS policy number specified on the egress SAP
E.Mac/IP	The MAC or IP filter policy ID applied to the egress SAP
A.Pol	The accounting policy ID assigned to the SAP
Adm	The administrative state of the SAP
Opr	The operational state of the SAP

sdp

Syntax **sdp** [*sdp-id* | **far-end** *ip-addr*ess] [**detail**]

Context show>service>id

Description This command displays information for the SDPs associated with the service.

If no optional parameters are specified, a summary of all associated SDPs is displayed.

Parameters *sdp-id* — Displays only information for the specified SDP ID

Values 1 to 17407

Default All SDPs

ip-address — Displays only SDPs matching with the specified far-end IP address

Default SDPs with any far-end IP address.

detail — Displays detailed SDP information

Output The following output is an example of service SDP information.

Sample Output

```
*A:ALA-48# show service id 700 sdp 2
===============================================================================
Service Destination Point (Sdp Id : 2)
===============================================================================
SdpId            Type IP address      Adm     Opr       I.Lbl       E.Lbl
-------------------------------------------------------------------------------
2:222            Spok 10.10.10.104    Up      Down      0           0
2:700            Mesh 10.10.10.104    Up      Down      0           0
===============================================================================
*A:ALA-48#
```

The following table describes show service-id SDP output fields.

Table 21 **Show Service-ID SDP Field Descriptions**

Label	Description
Sdp Id	The SDP identifier
Type	Indicates whether the SDP is a spoke or a mesh
Split Horizon Group	The name of the split horizon group that the SAP belongs
VC Type	The VC type: ether, vlan, or vpls
VC Tag	The explicit dot1Q value used when encapsulating to the SDP far end
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP
Admin Path MTU	The operating path MTU of the SDP is equal to the admin path MTU (when one is set) or the dynamically computed tunnel MTU, when no admin path MTU is set (the default case)

Table 21 **Show Service-ID SDP Field Descriptions (Continued)**

Label	Description (Continued)
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented
Far End	The IP address of the remote end of the GRE or MPLS tunnel defined by this SDP
Delivery	The type of delivery used by the SDP: GRE or MPLS
Admin State	The administrative state of this SDP
Oper State	The current status of the KeepAlive protocol
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP
Egress Label	The label used by this device to send packets to the far-end device in this service by the SDP
Last Changed	The date and time of the most recent change to the SDP
Signaling	The signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP
Admin State	The administrative state of the keepalive process
Oper State	The operational state of the keepalive process
Hello Time	The frequency that SDP echo request messages are transmitted
Max Drop Count	The maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault
Hello Msg Len	The length of the transmitted SDP echo request messages
Hold Down Time	The amount of time to wait before the keepalive operating status is eligible to enter the alive state
I. Fwd. Pkts.	The number of forwarded ingress packets
I. Dro. Pkts	The number of dropped ingress packets
E. Fwd. Pkts.	The number of forwarded egress packets
E. Fwd. Octets	The number of forwarded egress octets

Table 21 **Show Service-ID SDP Field Descriptions (Continued)**

Label	Description (Continued)
Associated LSP List	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the far end field. If the SDP type is GRE, then the following message displays: SDP delivery mechanism is not MPLS

sdp-using

Syntax **sdp-using** [*sdp-id*[:*vc-id*] | **far-end** *ip-address*]

Context show>service

Description This command displays services using SDP or far-end address options.

Parameters *sdp-id* — Displays only services bound to the specified SDP ID.

Values 1 to 17407

vc-id — Displays the virtual circuit identifier.

Values 1 to 4294967295

ip-address — Displays only services matching with the specified far-end IP address.

Default Services with any far-end IP address

Output The following output is an example of services using specified SDP or far-end address information.

Sample Output

```
A:ALA-1# show service sdp-using 300
===============================================================================
Service Destination Point (Sdp Id : 300)
===============================================================================
SvcId       SdpId               Type Far End        Opr State I.Label  E.Label
-------------------------------------------------------------------------------
1           300:1               Mesh 10.0.0.13      Up        131071   131071
2           300:2               Spok 10.0.0.13      Up        131070   131070
100         300:100             Mesh 10.0.0.13      Up        131069   131069
101         300:101             Mesh 10.0.0.13      Up        131068   131068
102         300:102             Mesh 10.0.0.13      Up        131067   131067
-------------------------------------------------------------------------------
Number of SDPs : 5
===============================================================================
A:ALA-1#
```

The following table describes **sdp-using** output fields.

Table 22 **Show Service SDP Using Field Descriptions**

Label	Description
Svc ID	The value ID entifying a service
Sdp ID	The SDP identifier
Type	The type of SDP: Spoke or Mesh
Far End	The far-end address of the SDP
Oper State	The operational state of the service
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP

service-using

Syntax **service-using** [**ies**] [**vpls**] [**vprn**] [**sdp** *sdp-id*]

Context show>service

Description This command displays the services matching certain usage properties.

If no parameters are specified, all services defined on the system are displayed.

Parameters **ies** — Displays matching IES instances.

vpls — Displays matching VPLS instances.

vprn — Displays matching VPRN instances (applies only to the 7750 SR).

sdp *sdp-id* — Displays only services bound to the specified SDP ID.

Values 1 to 17407

Default Services bound to any SDP ID

Output The following output is an example of service information of service types or SDPs.

Sample Output

```
*A:ALA-48>show>service# service-using vpls
===============================================================================
Services [vpls]
===============================================================================
ServiceId    Type      Adm    Opr        CustomerId        Last Mgmt Change
-------------------------------------------------------------------------------
700          VPLS      Up     Down       7                 04/11/2007 09:36:36
725          VPLS      Down   Down       7                 04/11/2007 09:36:36
```

```
740           VPLS      Down    Down       1                  04/11/2007 09:36:36
750           VPLS      Down    Down       7                  04/11/2007 09:36:36
1730          VPLS      Down    Down       1730               04/11/2007 09:36:36
9000          VPLS      Up      Down       6                  04/11/2007 09:36:36
90001         VPLS      Up      Down       6                  04/11/2007 09:36:36
-------------------------------------------------------------------------------
Matching Services : 7
===============================================================================
*A:ALA-48>show>service#
```

The following table describes show **service-using** output fields.

Table 23 **Show Service Service-Using Field Descriptions**

Label	Description
Service Id	The value that identifies a service
Type	The service type configured for the service ID
Adm	The administrative state of the service
Opr	The operational state of the service
CustomerID	The value that identifies a specific customer
Last Mgmt Change	The date and time of the most recent management-initiated change to this service

diameter

Syntax **diameter**

Context show>service>active-subscribers>hierarchy

Description This command displays active subscriber diameter information.

radius-acct

Syntax **radius-acct**

Context show>service>active-subscribers>hierarchy

Description This command provides information about subscriber RADIUS accounting sessions including: the accounting policy name, the accounting mode, the interim intervals, the accounting session ID, and the multi-session ID. If there are no RADIUS accounting sessions enabled for subscribers, this command still shows the session ID that can be used by CoA.

PD hosts modeled as routes are not considered hosts, and therefore, are not shown in the output of this command.

This command only shows a numeric RADIUS session ID even if the RADIUS accounting policy is configured using the descriptive format.

Output The following output is an example of subscriber RADIUS accounting session information.

Sample Output

```
A:BNG # show  service active-subscribers hierarchy radius-acct
===============================================================================
Active Subscribers Hierarchy (accounting information)
===============================================================================
-- home-1 (sub-prof-1)
   |   Acct-policy       : host-01
   |       Mode          : Queue Instance (Start/Stop), Host (Interim)
   |       Cfg. Interval : 5 min
   |
   +-- sap:1/1/20:841 - sla:sla-prof-1
       |   |
       |   +-- Queue Instance Acct - host-01
       |       session-id          : D854FF0000001E57507405
       |       multi-session-id    : D854FF0000001E57507405
       |
       +-- IPOE-Session - mac:00:00:10:10:12:12
           |   |
           |   +-- no radius acct
           |       session-id      : D854FF0000001C57507405
           |
           |-- 192.168.0.15
           |   |
           |   +-- Host Acct - host-01
           |       session-id      : D854FF000019C65750858C
           |       multi-session-id: D854FF0000001E57507405
           |
           +-- 2001:1000::1/128
               |
               +-- Host Acct - host-01
                   session-id      : D854FF0000001D57507405
                   multi-session-id: D854FF0000001E57507405
```

pcc-rule

Syntax **pcc rule** [**subscriber** *sub-ident-string*]

Context show>service>active-subscribers

Description This command specifies the pcc-rule.

Parameters *sub-ident-string* — Specifies the sub-ident-string, up to 32 characters.

all

Syntax **all**

Context show>service>id

Description This command displays detailed information for all aspects of the service.

Output Show All Service-ID Output

Sample Output

```
*A:PE-6# show service id 1 all
===============================================================================
Service Detailed Information
===============================================================================
Service Id        : 1                   Vpn Id            : 0
Service Type      : VPLS
Name              : 1
Description       : (Not Specified)
Customer Id       : 1                   Creation Origin   : manual
Last Status Change: 05/08/2018 09:40:32
Last Mgmt Change  : 05/08/2018 09:40:24
Etree Mode        : Disabled
Admin State       : Up                  Oper State        : Up
MTU               : 1514
SAP Count         : 1                   SDP Bind Count    : 1
Snd Flush on Fail : Disabled            Host Conn Verify  : Disabled
SHCV pol IPv4     : None
Propagate MacFlush: Disabled            Per Svc Hashing   : Disabled
Allow IP Intf Bind: Disabled
Fwd-IPv4-Mcast-To*: Disabled            Fwd-IPv6-Mcast-To*: Disabled
Mcast IPv6 scope  : mac-based
Def. Gateway IP   : None
Def. Gateway MAC  : None
Temp Flood Time   : Disabled            Temp Flood        : Inactive
Temp Flood Chg Cnt: 0
SPI load-balance  : Disabled
TEID load-balance : Disabled
Src Tep IP        : N/A
Vxlan ECMP        : Disabled
VSD Domain        : <none>
-------------------------------------------------------------------------------
BGP Information
-------------------------------------------------------------------------------
Split Horizon Group specifics
-------------------------------------------------------------------------------
ETH-CFM service specifics
-------------------------------------------------------------------------------
Tunnel Faults     : ignore
-------------------------------------------------------------------------------
Service Destination Points(SDPs)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Sdp Id 65:1  -(192.0.2.5)
-------------------------------------------------------------------------------
Description     : (Not Specified)
SDP Id            : 65:1                  Type              : Spoke
Spoke Descr     : (Not Specified)
Split Horiz Grp   : (Not Specified)
Etree Root Leaf Tag: Disabled             Etree Leaf AC     : Disabled
VC Type           : Ether                 VC Tag            : n/a
```

```
Admin Path MTU     : 0                        Oper Path MTU     : 8974
Delivery           : MPLS
Far End            : 192.0.2.5                Tunnel Far End    :
Oper Tunnel Far End: 192.0.2.5
LSP Types          : LDP
Hash Label         : Disabled                 Hash Lbl Sig Cap  : Disabled
Oper Hash Label    : Disabled
Entropy Label      : Disabled

Admin State        : Up                       Oper State        : Down
MinReqd SdpOperMTU : 1514
Acct. Pol          : None                     Collect Stats     : Disabled
Ingress Label      : 524279                   Egress Label      : None
Ingr Mac Fltr-Id   : n/a                      Egr Mac Fltr-Id   : n/a
Ingr IP Fltr-Id    : n/a                      Egr IP Fltr-Id    : n/a
Ingr IPv6 Fltr-Id  : n/a                      Egr IPv6 Fltr-Id  : n/a
Admin ControlWord  : Not Preferred            Oper ControlWord  : False
BFD Template       : None
BFD-Enabled        : no                       BFD-Encap         : ipv4
Last Status Change : 05/07/2018 16:50:10      Signaling         : TLDP
Last Mgmt Change   : 05/08/2018 09:40:24
Endpoint           : N/A                      Precedence        : 4
PW Status Sig      : Enabled
Force Vlan-Vc      : Disabled                 Force Qinq-Vc     : none
Class Fwding State : Down
Flags              : NoEgrVCLabel
Time to RetryReset : never                    Retries Left      : 3
Mac Move           : Blockable                Blockable Level   : Tertiary
Local Pw Bits      : None
Peer Pw Bits       : None
Peer Fault Ip      : None
Peer Vccv CV Bits  : None
Peer Vccv CC Bits  : None

Application Profile: None
Max Nbr of MAC Addr: No Limit                 Total MAC Addr    : 0
Learned MAC Addr   : 0                        Static MAC Addr   : 0
OAM MAC Addr       : 0                        DHCP MAC Addr     : 0
Host MAC Addr      : 0                        Intf MAC Addr     : 0
SPB MAC Addr       : 0                        Cond MAC Addr     : 0
BGP EVPN Addr      : 0                        EVPN Static Addr  : 0

MAC Learning       : Enabled                  Discard Unkwn Srce: Disabled
MAC Aging          : Enabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
MAC Pinning        : Disabled
Ignore Standby Sig : False                    Block On Mesh Fail: False
Oper Group         : (none)                   Monitor Oper Grp  : (none)
Auto Learn Mac Prot: Disabled
RestMacProtSrc Act : none

Ingress Qos Policy : (none)                   Egress Qos Policy : (none)
Ingress FP QGrp    : (none)                   Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                   Egr Port QGrp Inst: (none)

KeepAlive Information :
Admin State        : Disabled                 Oper State        : Disabled
Hello Time         : 10                       Hello Msg Len     : 0
```

```
Max Drop Count     : 3                          Hold Down Time    : 10

Statistics           :
I. Fwd. Pkts.      : 0                          I. Dro. Pkts.     : 0
I. Fwd. Octs.      : 0                          I. Dro. Octs.     : 0
E. Fwd. Pkts.      : 0                          E. Fwd. Octets    : 0
-------------------------------------------------------------------------------
Control Channel Status
-------------------------------------------------------------------------------
PW Status          : disabled                   Refresh Timer     : <none>
Peer Status Expire : false
Request Timer      : <none>
Acknowledgement    : false

MCAC Policy Name   :
MCAC Max Unconst BW: no limit                   MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                          MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                          MCAC Avail Opnl BW: unlimited
MCAC If-Policy Name:
-------------------------------------------------------------------------------
ETH-CFM SDP-Bind specifics
-------------------------------------------------------------------------------
V-MEP Filtering    : Disabled
Squelch Levels     : None
Collect Lmm Stats  : Disabled
LMM FC Stats       : None
LMM FC In Prof     : None
-------------------------------------------------------------------------------
LDP Information :
-------------------------------------------------------------------------------
LDP LSP Id         : 65541
-------------------------------------------------------------------------------
RSVP/Static LSPs
-------------------------------------------------------------------------------
Associated LSP List :
No LSPs Associated
-------------------------------------------------------------------------------
Class-based forwarding :
-------------------------------------------------------------------------------
Class forwarding   : Disabled                   EnforceDSTELspFc  : Disabled
Default LSP        : Uknwn
Multicast LSP      : None
===============================================================================
FC Mapping Table
===============================================================================
FC Name             LSP Name
-------------------------------------------------------------------------------
No FC Mappings
-------------------------------------------------------------------------------
Segment Routing
-------------------------------------------------------------------------------
ISIS               : disabled
OSPF               : disabled
TE-LSP             : disabled
-------------------------------------------------------------------------------
Stp Service Destination Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                         Stp Oper State    : Down
Core Connectivity  : Down
```

```
Port Role          : N/A                       Port State        : Discarding
Port Number        : 0                         Port Priority     : 128
Port Path Cost     : 10                        Auto Edge         : Enabled
Admin Edge         : Disabled                  Oper Edge         : N/A
Link Type          : Pt-pt                     BPDU Encap        : Dot1d
Root Guard         : Disabled                  Active Protocol   : N/A
Last BPDU from     : N/A
Designated Bridge  : N/A                       Designated Port Id: 0

Fwd Transitions    : 0                         Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                         Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                         TCN BPDUs tx      : 0
TC bit BPDUs rcvd  : 0                         TC bit BPDUs tx   : 0
RST BPDUs rcvd     : 0                         RST BPDUs tx      : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
-------------------------------------------------------------------------------
Service Access Points
-------------------------------------------------------------------------------
SAP 1/1/c1/1:1
-------------------------------------------------------------------------------
Service Id         : 1
SAP                : 1/1/c1/1:1                Encap             : q-tag
Description        : (Not Specified)
Admin State        : Up                        Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 05/07/2018 16:50:10
Last Mgmt Change   : 05/08/2018 09:40:32
Sub Type           : regular
Dot1Q Ethertype    : 0x8100                    QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Etree Root Leaf Tag: Disabled                  Etree Leaf Tag    : 0
Etree Leaf AC      : Disabled
Max Nbr of MAC Addr: No Limit                  Total MAC Addr    : 0
Learned MAC Addr   : 0                         Static MAC Addr   : 0
OAM MAC Addr       : 0                         DHCP MAC Addr     : 0
Host MAC Addr      : 0                         Intf MAC Addr     : 0
SPB MAC Addr       : 0                         Cond MAC Addr     : 0
BGP EVPN Addr      : 0                         EVPN Static Addr  : 0
Admin MTU          : 9000                      Oper MTU          : 9000
Ingr IP Fltr-Id    : n/a                       Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                       Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                       Egr IPv6 Fltr-Id  : n/a
qinq-pbit-marking  : both
Egr Agg Rate Limit : max
Q Frame-Based Acct : Disabled                  Limit Unused BW   : Disabled
ARP Reply Agent    : Disabled                  Host Conn Verify  : Disabled
SHCV pol IPv4      : None
Mac Learning       : Enabled                   Discard Unkwn Srce: Disabled
Mac Aging          : Enabled                   Mac Pinning       : Disabled
BPDU Translation   : Disabled
L2PT Termination   : Disabled
Vlan-translation   : None
Qinq-vlan-                                     Qinq-vlan-
translation        : None                      translation Ids   : None

Acct. Pol          : None                      Collect Stats     : Disabled
```

```
Anti Spoofing      : None                      Dynamic Hosts     : Enabled
Avl Static Hosts   : 0                         Tot Static Hosts  : 0
Calling-Station-Id : n/a

Oper Group         : (none)                    Monitor Oper Grp  : (none)
Host Lockout Plcy  : n/a
Lag Link Map Prof  : (none)
Cflowd             : Disabled
Bandwidth          : Not-Applicable
Oper DCpu Prot Pol*: _default-access-policy
MCAC Policy Name   :                           MCAC Const Adm St : Enable
MCAC Max Unconst BW: no limit                  MCAC Max Mand BW  : no limit
MCAC In use Mand BW: 0                         MCAC Avail Mand BW: unlimited
MCAC In use Opnl BW: 0                         MCAC Avail Opnl BW: unlimited
Use LAG port weight: no
MCAC If-Policy Name:
Restr MacUnpr Dst  : Disabled
Auto Learn Mac Prot: Disabled
RestMacProtSrc Act : none
Time to RetryReset : never                     Retries Left      : 3
Mac Move           : Blockable                 Blockable Level   : Tertiary
Auth Policy        : None
DestMac Rewrite    : Disabled
SendBvplsEvpnFlush : Enabled
-------------------------------------------------------------------------------
ETH-CFM SAP specifics
-------------------------------------------------------------------------------
Tunnel Faults      : n/a                       AIS               : Disabled
MC Prop-Hold-Timer : n/a                       V-MEP Filtering   : Disabled
Squelch Levels     : None
Collect Lmm Stats  : Disabled
LMM FC Stats       : None
LMM FC In Prof     : None
-------------------------------------------------------------------------------
Stp Service Access Point specifics
-------------------------------------------------------------------------------
Stp Admin State    : Up                        Stp Oper State    : Down
Core Connectivity  : Down
Port Role          : N/A                       Port State        : Forwarding
Port Number        : N/A                       Port Priority     : 128
Port Path Cost     : 10                        Auto Edge         : Enabled
Admin Edge         : Disabled                  Oper Edge         : N/A
Link Type          : Pt-pt                     BPDU Encap        : Dot1d
Root Guard         : Disabled                  Active Protocol   : N/A
Last BPDU from     : N/A
CIST Desig Bridge  : N/A                       Designated Port   : N/A

Forward transitions: 0                         Bad BPDUs rcvd    : 0
Cfg BPDUs rcvd     : 0                         Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                         TCN BPDUs tx      : 0
TC bit BPDUs rcvd  : 0                         TC bit BPDUs tx   : 0
RST BPDUs rcvd     : 0                         RST BPDUs tx      : 0
MST BPDUs rcvd     : 0                         MST BPDUs tx      : 0
-------------------------------------------------------------------------------
ARP host
-------------------------------------------------------------------------------
Admin State        : outOfService
Host Limit         : 1                         Min Auth Interval : 15 minutes
```

```
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                       Egress qos-policy : 1
Ingress FP QGrp    : (none)                  Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                  Egr Port QGrp Inst: (none)
Shared Q plcy      : n/a                     Multipoint shared : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
I. Policer Ctl Pol : (Not Specified)
E. Policer Ctl Pol : (Not Specified)
I. QGrp Redir. List: (Not Specified)
E. QGrp Redir. List: (Not Specified)
-------------------------------------------------------------------------------
DHCP
-------------------------------------------------------------------------------
Description        : (Not Specified)
Admin State        : Down                    Lease Populate    : 0
DHCP Snooping      : Down                    Action            : Keep

Proxy Admin State  : Down
Proxy Lease Time   : N/A
Emul. Server Addr  : Not Configured
-------------------------------------------------------------------------------
DHCP6
-------------------------------------------------------------------------------
Description        : (Not Specified)
Admin State        : Down
DHCP Snooping      : Down
Interface-Id       : AsciiTuple              Remote-Id         : Disabled
-------------------------------------------------------------------------------
Subscriber Management
-------------------------------------------------------------------------------
Admin State        : Down                    MAC DA Hashing    : False
Def Sub-Id         : None
Def Sub-Profile    : None
Def SLA-Profile    : None
Def Inter-Dest-Id  : None
Sub-Ident-Policy   : None

Subscriber Limit   : 1
Single-Sub-Parameters
Prof Traffic Only : False
Non-Sub-Traffic   : N/A

Static host management
MAC learn options  : N/A

-------------------------------------------------------------------------------
Sap Aggregate Stats
-------------------------------------------------------------------------------
                       Packets                 Octets
Ingress
Aggregate Offered     : 0                      0
Aggregate Forwarded   : 0                      0
Aggregate Dropped     : 0                      0

Egress
```

```
Aggregate Forwarded   : 0                       0
Aggregate Dropped     : 0                       0
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                 Octets
CPM Ingress           : 0                       0

Forwarding Engine Stats
Dropped               : 0                       0
Received Valid        : 0                       0
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Off. Uncolor          : 0                       0
Off. Managed          : 0                       0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Queueing Stats(Egress QoS Policy 1)
Dro. In/InplusProf    : 0                       0
Dro. Out/ExcProf      : 0                       0
For. In/InplusProf    : 0                       0
For. Out/ExcProf      : 0                       0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Queue 11 (Multipoint) (Priority)
Off. Combined         : 0                       0
Off. Managed          : 0                       0
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Egress Queue 1
For. In/InplusProf    : 0                       0
For. Out/ExcProf      : 0                       0
Dro. In/InplusProf    : 0                       0
Dro. Out/ExcProf      : 0                       0
* indicates that the corresponding row element may have been truncated.

-------------------------------------------------------------------------------
VPLS Spanning Tree Information
```

```
-------------------------------------------------------------------------------
VPLS oper state    : Up                        Core Connectivity : Down
Stp Admin State    : Down                      Stp Oper State    : Down
Mode               : Rstp                      Vcp Active Prot.  : N/A

Bridge Id          : 80:00.d8:4a:ff:00:00:00  Bridge Instance Id: 0
Bridge Priority    : 32768                     Tx Hold Count     : 6
Topology Change    : Inactive                  Bridge Hello Time : 2
Last Top. Change   : 0d 00:00:00               Bridge Max Age    : 20
Top. Change Count  : 0                         Bridge Fwd Delay  : 15
MST region revision: 0                         Bridge max hops   : 20
MST region name    :

Root Bridge        : N/A
Primary Bridge     : N/A

Root Path Cost     : 0                         Root Forward Delay: 0
Rcvd Hello Time    : 0                         Root Max Age      : 0
Root Priority      : 0                         Root Port         : N/A
-------------------------------------------------------------------------------
Forwarding Database specifics
-------------------------------------------------------------------------------
Service Id        : 1                   Mac Move          : Disabled
Primary Factor    : 3                   Secondary Factor  : 2
Mac Move Rate     : 2                   Mac Move Timeout  : 10
Mac Move Retries  : 3
Table Size        : 250                 Allocated Count   : 0
Total In Use      : 0
Learned Count     : 0                   Static Count      : 0
OAM MAC Count     : 0                   DHCP MAC Count    : 0
Host MAC Count    : 0                   Intf MAC Count    : 0
Spb Count         : 0                   Cond MAC Count    : 0
BGP EVPN Count    : 0                   EVPN Static Cnt   : 0
EVPN Dup Det Cnt  : 0
Remote Age        : 900                 Local Age         : 300
High Watermark    : 95%                 Low Watermark     : 90%
Mac Learning      : Enabled             Discard Unknown   : Disabled
Mac Aging         : Enabled             Relearn Only      : False
Mac Subnet Len    : 48
Sel Learned FDB   : Disabled
-------------------------------------------------------------------------------
IGMP Snooping Base info
-------------------------------------------------------------------------------
Admin State : Down
Querier     : No querier found
-------------------------------------------------------------------------------
Port                     Oper MRtr Pim  Send Max   Max  Max   MVR       Num
Id                       Stat Port Port Qrys Grps  Srcs Grp   From-VPLS Grps
                                                       Srcs
-------------------------------------------------------------------------------
sap:1/1/c1/1:1           Up   No   No   No   None  None None  Local     0
sdp:65:1                 Down No   No   No   None  None None  N/A       0
-------------------------------------------------------------------------------
MLD Snooping Base info
-------------------------------------------------------------------------------
Admin State : Down
Querier     : No querier found
-------------------------------------------------------------------------------
Port                    Oper    MRtr  Send      Max Num   MVR        Num
```

```
Id                        State   Port  Queries   Groups    From-VPLS  Groups
-------------------------------------------------------------------------------
sap:1/1/c1/1:1            Up      No    Disabled  No Limit  Local      0
sdp:65:1                  Down    No    Disabled  No Limit  N/A        0
-------------------------------------------------------------------------------
DHCP Summary, service 1
-------------------------------------------------------------------------------
Sap/Sdp                 Snoop  Used/       Arp Reply   Info     Admin
                               Provided    Agent       Option   State
-------------------------------------------------------------------------------
sap:1/1/c1/1:1          No     0/0         No          Keep     Down
sdp:65:1                No     N/A         N/A         N/A      N/A
-------------------------------------------------------------------------------
Number of Entries : 2
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
ARP host Summary, service 1
-------------------------------------------------------------------------------
Sap                     Used        Provided    Admin State
-------------------------------------------------------------------------------
sap:1/1/c1/1:1          0           1           outOfService
-------------------------------------------------------------------------------
Number of SAPs : 1     0
-------------------------------------------------------------------------------
===============================================================================
Service VPLS Group Information
===============================================================================
===============================================================================
No vxlan information for 1
-------------------------------------------------------------------------------
Service Endpoints
-------------------------------------------------------------------------------
No Endpoints found.
-------------------------------------------------------------------------------
===============================================================================
VPLS Sites
===============================================================================
Site                Site-Id   Dest                Mesh-SDP  Admin   Oper  Fwdr
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

The following table describes the show all service-id command output fields:

Table 24 **Show All Service-ID Output Field Descriptions**

Label	Description
Service Id	The value that identifies a service
VPN Id	The number that identifies the VPN
Service Type	The type of service
SDP Id	The SDP identifier

Table 24 **Show All Service-ID Output Field Descriptions (Continued)**

Label	Description (Continued)
Description	The text string describing general information about the service
Customer Id	The customer identifier
Last Mgmt Change	The date and time of the most recent management-initiated change to this customer
SAP Count	The number of SAPs specified for this service
SDP Bind Count	The number of SDPs bound to this service
Split Horizon Group	The name of the split horizon group for this service
Description	The text string describing the split horizon group
Last Changed	The date and time of the most recent management-initiated change to this split horizon group
SDP Id	The SDP identifier
Type	The service SDP binding type, spoke or mesh
Admin Path MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Delivery	The type of delivery used by the SDP: GRE or MPLS
Admin State	The administrative state of this SDP
Oper State	The operational status of the KeepAlive protocol.
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP
Egress Label	The label used by this device to send packets to the far-end device in this service by this SDP
Ingress Filter	The ID of the ingress filter policy
Egress Filter	The ID of the egress filter policy
Far End	The IP address of the remote end of the GRE or MPLS tunnel defined by this SDP
Last Changed	The date and time of the most recent change to this customer

Table 24 **Show All Service-ID Output Field Descriptions (Continued)**

Label	Description (Continued)
Signaling	The signaling protocol used to obtain the ingress and egress labels used in frames transmitted and received on this SDP.
Admin State	The operating status of the SDP
Oper State	The operational state of the SDP
Hello Time	The frequency that SDP echo request messages are transmitted on this SDP
Hello Msg Len	The length of the SDP echo request messages transmitted on this SDP.
Max Drop Count	The maximum number of consecutive SDP echo request messages that can be unacknowledged before the keepalive protocol reports a fault
Hold Down Time	The amount of time to wait before the keepalive operating status is eligible to enter the alive state.
SDP Delivery Mechanism	When the SDP type is MPLS, a list of LSPs used to reach the far-end router displays. All the LSPs in the list must terminate at the IP address specified in the Far End field. If the SDP type is GRE, then the following message displays: SDP Delivery Mechanism is not MPLS
Number of SDPs	The total number SDPs applied to this service ID
Service Access Points	
Service Id	The value that identifies a service
Port Id	The ID of the access port where this SAP is defined
Description	The generic information about the SAP
Encap Value	The value of the label used to identify this SAP on the access port.
Admin State	The desired state of the SAP
Oper State	The operating state of the SAP
Last Changed	The date and time of the last change
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented

Table 24 **Show All Service-ID Output Field Descriptions (Continued)**

Label	Description (Continued)
Ingress qos-policy	The SAP ingress QoS policy ID
Egress qos-policy	The SAP egress QoS policy ID
Ingress Filter-Id	The SAP ingress filter policy ID
Egress Filter-Id	The SAP egress filter policy ID
Multi Svc Site	The multi-service site in which the SAP is a member
Ingress sched-policy	The ingress QoS scheduler for the SAP
Egress sched-policy	The egress QoS scheduler for the SAP
Acct. Pol	The accounting policy applied to the SAP
Collect Stats	Accounting statistics collected on the SAP
Dropped	The number of packets or octets dropped
Offered Hi Priority	The number of high priority packets, as determined by the SAP ingress QoS policy
Offered Low Priority	The number of low priority packets, as determined by the SAP ingress QoS policy
Forwarded In Profile	The number of in-profile packets or octets (rate below CIR) forwarded
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded
Dropped In Profile	The number of in-profile packets or octets discarded
Dropped Out Profile	The number of out-of-profile packets or octets discarded
Forwarded In Profile	The number of in-profile packets or octets (rate below CIR) forwarded
Forwarded Out Profile	The number of out-of-profile packets or octets (rate above CIR) forwarded
Ingress Queue 1	The index of the ingress QoS queue of this SAP
High priority offered	The packets or octets count of the high priority traffic for the SAP
High priority dropped	The number of high priority traffic packets or octets dropped
Low priority offered	The packets or octets count of the low priority traffic
Low priority dropped	The number of low priority traffic packets or octets dropped

Table 24 **Show All Service-ID Output Field Descriptions (Continued)**

Label	Description (Continued)
In profile forwarded	The number of in-profile packets or octets (rate below CIR) forwarded
Out profile forwarded	The number of out-of-profile octets (rate above CIR) forwarded
Egress Queue 1	The index of the egress QoS queue of the SAP
In profile forwarded	The number of in-profile packets or octets (rate below CIR) forwarded
In profile dropped	The number of in-profile packets or octets dropped for the SAP
Out profile forwarded	The number of out-of-profile packets or octets (rate above CIR) forwarded
Out profile dropped	The number of out-of-profile packets or octets discarded
State	The DHCP relay state on this SAP
Info Option	Option 82 processing state on this SAP
Action	Option 82 processing state SAP or interface: keep, replace or drop
Circuit ID	If index is inserted in Circuit ID sub-option of Option 82
Remote ID	The far-end MAC address inserted in remote ID sub-option of Option 82
Managed by Service	The service-id of the management VPLS managing this SAP
Managed by SAP	The SAP ID inside the management VPLS managing this SAP
Prune state	The STP state inherited from the management VPLS
Managed by Service	The service ID of the management VPLS managing this spoke SDP
Managed by Spoke	The SAP ID inside the management VPLS managing this spoke SDP
Prune state	The STP state inherited from the management VPLS

arp-host

Syntax **arp-host** [**wholesaler** *service-id*] [**sap** *sap-id* | **interface** *interface-name* | **ip-address** *ip-address*[/*mask*] | **mac** *ieee-address* | {[**port** *port-id*] [**no-inter-dest-id** | **inter-dest-id** *inter-dest-id*]}] [**detail**]

arp-host statistics [**sap** *sap-id* | **interface** *interface-name*]

arp-host summary [**interface** *interface-name*]

Context show>service>id

Description This command displays ARP host related information.

Output The following output is an example ARP host information for specified service IDs.

Sample Output

```
*A:Dut-C# show service id 2 arp-host
===============================================================================
ARP host table, service 2
===============================================================================
IP Address      Mac Address       Sap Id              Remaining           MC
                                                      Time                Stdby
-------------------------------------------------------------------------------
128.128.1.2     00:80:00:00:00:01 2/1/5:2             00h04m41s
128.128.1.3     00:80:00:00:00:02 2/1/5:2             00h04m42s
128.128.1.4     00:80:00:00:00:03 2/1/5:2             00h04m43s
128.128.1.5     00:80:00:00:00:04 2/1/5:2             00h04m44s
128.128.1.6     00:80:00:00:00:05 2/1/5:2             00h04m45s
128.128.1.7     00:80:00:00:00:06 2/1/5:2             00h04m46s
128.128.1.8     00:80:00:00:00:07 2/1/5:2             00h04m47s
128.128.1.9     00:80:00:00:00:08 2/1/5:2             00h04m48s
128.128.1.10    00:80:00:00:00:09 2/1/5:2             00h04m49s
128.128.1.11    00:80:00:00:00:0a 2/1/5:2             00h04m50s
-------------------------------------------------------------------------------
Number of ARP hosts : 10
===============================================================================
*A:Dut-C#


*A:Dut-C# show service id 2 arp-host ip-address 128.128.1.2 detail
===============================================================================
ARP hosts for service 2
===============================================================================
Service ID           : 2
IP Address           : 128.128.1.2
MAC Address          : 00:80:00:00:00:01
SAP                  : 2/1/5:2
Remaining Time       : 00h04m58s

Sub-Ident            : "alu_1_2"
Sub-Profile-String   : ""
SLA-Profile-String   : ""
App-Profile-String   : ""
ARP host ANCP-String : ""
ARP host Int Dest Id : ""

RADIUS-User-Name     : "128.128.1.2"

Session Timeout (s)  : 301
Start Time           : 02/09/2009 16:35:07
Last Auth            : 02/09/2009 16:36:34
Last Refresh         : 02/09/2009 16:36:38
Persistence Key      : N/A
-------------------------------------------------------------------------------
```

```
Number of ARP hosts : 1
===============================================================================
*A:Dut-C#


*A:Dut-C# show service id 2 arp-host statistics
===============================================================================
ARP host statistics
===============================================================================
Num Active Hosts             : 20
Received Triggers            : 70
Ignored Triggers             : 10
Ignored Triggers (overload)  : 0
SHCV Checks Forced           : 0
Hosts Created                : 20
Hosts Updated                : 40
Hosts Deleted                : 0
Authentication Requests Sent : 40
===============================================================================
*A:Dut-C#


*A:Dut-C# show service id 2 arp-host summary
============================================================
ARP host Summary, service 2
============================================================
Sap                    Used        Provided    Admin State
------------------------------------------------------------
sap:2/1/5:2            20          8000        inService
------------------------------------------------------------
Number of SAPs : 1
------------------------------------------------------------
============================================================
*A:Dut-C#
```

fdb

Syntax **fdb** [**sap** *sap-id*] | [**sdp** *sdp-id*] | [**mac** *ieee-address*] | [**detail**]

Context show>service>id
show>service>fdb-mac

Description This command displays FDB entry for a given MAC address.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition

sdp-id — Displays the SDP identifier

Values 1 to 17407

ieee-address — Specifies to display information pertaining to the MAC address

detail — Displays detailed information

Output The following output is an example of service FDB MAC information.

Sample Output

```
*A:ALA-48>show>service>id# fdb mac
===============================================================================
Service Forwarding Database
===============================================================================
ServId    MAC               Source-Identifier     Type/Age  Last Change
-------------------------------------------------------------------------------
6         00:aa:00:00:00:00 sap:lag-2             L/0       04/11/2007
15:04:31
6         00:aa:00:00:00:01 sap:lag-2             L/0       04/11/2007
15:04:31
6         00:aa:00:00:00:02 sap:lag-2             L/0       04/11/2007
15:04:31
6         00:aa:00:00:00:03 sap:lag-2             L/0       04/11/2007
15:04:31
6         00:aa:00:00:00:04 sap:lag-2             L/0       04/11/2007
15:04:31
10        12:12:12:12:12:12 sap:1/1/1:100         S         04/11/2007
10:03:29
===============================================================================
*A:ALA-48>show>service>id#
```

labels

Syntax **labels**

Context show>service>id

Description This command displays the labels being used by the service.

Output The following output is an example of service label information.

Sample Output

```
A:ALA-12# show service id 1 labels
===============================================================================
Martini Service Labels
===============================================================================
Svc Id     Sdp Id             Type I.Lbl        E.Lbl
-------------------------------------------------------------------------------
1          10:1               Mesh 0            0
1          20:1               Mesh 0            0
1          30:1               Mesh 0            0
1          40:1               Mesh 130081       131061
1          60:1               Mesh 131019       131016
1          100:1              Mesh 0            0
-------------------------------------------------------------------------------
Number of Bound SDPs : 6
-------------------------------------------------------------------------------
A:ALA-12#
```

The following table describes show service-id labels output fields:

Table 25 **Show Service-ID Labels Field Descriptions**

Label	Description
Svc Id	The service identifier
Sdp Id	The SDP identifier
Type	The SDP type, spoke or mesh
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP
E. Lbl	The VC label used by this device to send packets to the far-end device in this service by the SDP

mfib

Syntax **mfib brief**
mfib [**group** *grp-address*]
mfib statistics [**group** *grp-address*]

Context show>service>id

Description This command displays the multicast FIB on the VPLS service.

Parameters *grp-ip-address* — Displays the multicast FIB for a specific multicast group address.

statistics — Displays statistics on the multicast FIB.

Output Show Output

The following table describes the command output fields:

Table 26 **Triple Play Mfib Output Field Descriptions**

Label	Description
Source Address	The IPv4 unicast source address
Group Address	The IPv4 multicast group address
SAP/SDP ID	The SAP or SDP to which the corresponding multicast stream is forwarded or blocked
Forwarding/Blocking	The corresponding multicast stream status, blocked or forwarded
Number of Entries	The number of entries in the MFIB
Forwarded Packets	The number of multicast packets forwarded for the corresponding source or group

Table 26 **Triple Play Mfib Output Field Descriptions (Continued)**

Label	Description (Continued)
Svc ID	The service to which the corresponding multicast stream is forwarded or blocked. Local means that the multicast stream is forwarded or blocked to a SAP or SDP local to the service.

Sample Output

The following is an example for the 7450 ESS:

```
A:ALA-1>show>service>id # mfib
===============================================================================
Multicast FIB, Service 10
===============================================================================
Source Address  Group Address   Sap/Sdp Id             Svc Id     Fwd/Blk
-------------------------------------------------------------------------------
*               239.1.1.1       sap:1/1/1:10           Local      Fwd
                                sap:1/1/2              20         Fwd
                                sap:1/1/4:100          20         Fwd
                                sap:1/1/4:200          20         Fwd
*               239.1.1.2       sap:1/1/4:200          20         Fwd
-------------------------------------------------------------------------------
Number of entries: 2
===============================================================================
A:ALA-1>
```

The following is an example for the 7750 SR:

```
A:ALA-1>show>service>id # mfib
===============================================================================
IGMP Snooping MFIB for service 10
===============================================================================
Source Address  Group Address   Sap/Sdp Id             Fwd/Blk
-------------------------------------------------------------------------------
*               239.0.0.1       sap:2/1/5:1            Fwd
*               239.0.0.7       sap:2/1/5:7            Fwd
-------------------------------------------------------------------------------
Number of entries: 2


*A:rbae_C# show service id 1 mfib statistics
===============================================================================
Multicast FIB Statistics, Service 1
===============================================================================
Source Address  Group Address       Matched Pkts          Matched Octets
-------------------------------------------------------------------------------
*               *                   838                   50280
10.11.0.0       239.0.0.0           8                     480
10.11.0.0       239.0.0.1           0                     0
*               * (MLD)             0                     0
*               33:33:00:00:00:01   2650                  159000
*               33:33:00:00:00:02   0                     0
-------------------------------------------------------------------------------
Number of entries: 6
===============================================================================
```

```
*A:rbae_C#
```

mld-snooping

Syntax **mld-snooping**

Context show>service>id

Description This command displays MLD snooping information.

all

Syntax **all**

Context show>service>id>mld-snooping

Description This command displays detailed information about MLD snooping.

Output The following output is an example of service MLD snooping information.

Sample Output

```
*A:rbae_C# show service id 1 mld-snooping all
===============================================================================
MLD Snooping info for service 1
-------------------------------------------------------------------------------
MLD Snooping Base info
-------------------------------------------------------------------------------
Admin State : Up
Querier     : FE80:db8:12 on SAP 2/1/5
-------------------------------------------------------------------------------
Sap/Sdp                  Oper    MRtr  Send      Max Num   MVR        Num
Id                       State   Port  Queries   Groups    From-VPLS  Groups
-------------------------------------------------------------------------------
sap:1/1/4                Up      No    Disabled  No Limit  Local      0
sap:2/1/5                Up      Yes   Disabled  No Limit  Local      0
sdp:31:1                 Up      No    Disabled  No Limit  N/A        0
sdp:36:1                 Up      No    Disabled  No Limit  N/A        0
-------------------------------------------------------------------------------
MLD Snooping Querier info
-------------------------------------------------------------------------------
Sap Id                  : 2/1/5
IP Address              : FE80:db8:12
Expires                 : 11s
Up Time                 : 0d 00:05:05
Version                 : 2

General Query Interval  : 10s
Query Response Interval : 1.0s
Robust Count            : 2
-------------------------------------------------------------------------------
MLD Snooping Multicast Routers
```

```
-------------------------------------------------------------------------------
MRouter
                 Sap/Sdp Id             Up Time        Expires   Version
-------------------------------------------------------------------------------
FE80::12
                 2/1/5                  0d 00:05:05    11s       2
-------------------------------------------------------------------------------
Number of mrouters: 1
-------------------------------------------------------------------------------
MLD Snooping Proxy-reporting DB
-------------------------------------------------------------------------------
Group Address
                  Mode     Up Time        Num Sources
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
MLD Snooping SAP 1/1/4 Port-DB
-------------------------------------------------------------------------------
Group Address
                 Mode    Type    From-VPLS  Up Time        Expires  Num   MC
                                                                     Src   Stdby
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
MLD Snooping SAP 2/1/5 Port-DB
-------------------------------------------------------------------------------
Group Address
                 Mode    Type    From-VPLS  Up Time        Expires  Num   MC
                                                                     Src   Stdby
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
MLD Snooping SDP 31:1 Port-DB
-------------------------------------------------------------------------------
Group Address
                 Mode    Type    From-VPLS  Up Time        Expires  Num Src
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
MLD Snooping SDP 36:1 Port-DB
-------------------------------------------------------------------------------
Group Address
                 Mode    Type    From-VPLS  Up Time        Expires  Num Src
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
MLD Snooping Static Source Groups
-------------------------------------------------------------------------------
MLD Snooping Statistics
-------------------------------------------------------------------------------
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         43            0             129
Group Queries           0             0             0
Group-Source Queries    0             0             0
V1 Reports              0             0             0
V2 Reports              4             35            0
V1 Done                 0             0             0
Unknown Type            0             N/A           0
```

```
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length              : 0
Bad MLD Checksum        : 0
Bad Encoding            : 0
No Router Alert         : 0
Zero Source IP          : 0
Wrong Version           : 0
Lcl-Scope Packets       : 0
Rsvd-Scope Packets      : 0

Send Query Cfg Drops    : 0
Import Policy Drops     : 0
Exceeded Max Num Groups : 0
MCAC Policy Drops       : 0
MCS Failures            : 0

MVR From VPLS Cfg Drops : 0
MVR To SAP Cfg Drops    : 0
-------------------------------------------------------------------------------
MLD Snooping Multicast VPLS Registration info
-------------------------------------------------------------------------------
MLD Snooping Admin State  : Up

MVR Admin State           : Down
MVR Policy                : None
-------------------------------------------------------------------------------
Local SAPs/SDPs
-------------------------------------------------------------------------------
Svc Id     Sap/Sdp                Oper      From       Num Local
           Id                     State     VPLS       Groups
-------------------------------------------------------------------------------
1          sap:1/1/4              Up        Local      0
1          sap:2/1/5              Up        Local      0
1          sdp:31:1               Up        N/A        0
1          sdp:36:1               Up        N/A        0
-------------------------------------------------------------------------------
MVR SAPs (from-vpls=1)
-------------------------------------------------------------------------------
Svc Id     Sap/Sdp                Oper      From       Num MVR
           Id                     State     VPLS       Groups
-------------------------------------------------------------------------------
No MVR SAPs found.
===============================================================================
*A:rbae_C#
```

base

Syntax **base**

Context show>service>id>mld-snooping

Description This command displays basic MLD snooping information.

mrouters

Syntax **mrouters** [**detail**]

Context show>service>id>mld-snooping

Description This command displays all multicast routers.

Output The following output is an example of service MLD snooping mrouter information.

Sample Output

```
*A:rbae_C# show service id 1 mld-snooping mrouters
===============================================================================
MLD Snooping Multicast Routers for service 1
===============================================================================
MRouter
                 Sap/Sdp Id             Up Time        Expires   Version
-------------------------------------------------------------------------------
FE80::12
                 2/1/5                  0d 00:05:55    11s       2
-------------------------------------------------------------------------------
Number of mrouters: 1
===============================================================================
*A:rbae_C#


*A:rbae_C#  show service id 1 mld-snooping mrouters detail
===============================================================================
MLD Snooping Multicast Routers for service 1
-------------------------------------------------------------------------------
MRouter FE80::12
-------------------------------------------------------------------------------
Sap Id                  : 2/1/5
Expires                 : 18s
Up Time                 : 0d 00:06:28
Version                 : 2
-------------------------------------------------------------------------------
Number of mrouters: 1
```

mvr

Syntax **mvr**

Context show>service>id>mld-snooping

Description This command displays multicast VPLS registration information.

port-db

Syntax **port-db sap** *sap-id*

port-db sap *sap-id* **detail**
port-db sap *sap-id* **group** *grp-ipv6-address*
port-db sdp *sdp-id*:*vc-id* [**detail**]
port-db sdp *sdp-id*:*vc-id* **group** *grp-ipv6-address*

Context show>service>id>mld-snooping

Description This command displays MLD snooping information related to a specific SAP.

Output The following output is an example of MLD snooping information for a specific SAP.

Sample Output

```
*A:rbae_C# show service id 1 mld-snooping port-db sap 1/1/4
===============================================================================
MLD Snooping SAP 1/1/4 Port-DB for service 1
===============================================================================
Group Address
                Mode    Type    From-VPLS  Up Time         Expires  Num   MC
                                                                    Src   Stdby
-------------------------------------------------------------------------------
FF04::1
                include dynamic local      0d 00:00:19    0s       1
FF04::2
                include dynamic local      0d 00:00:18    0s       1
-------------------------------------------------------------------------------
Number of groups: 2
===============================================================================
*A:rbae_C#

*A:rbae_C# show service id 1 mld-snooping port-db sap 1/1/4 detail
===============================================================================
MLD Snooping SAP 1/1/4 Port-DB for service 1
-------------------------------------------------------------------------------
MLD Group FF04::1
-------------------------------------------------------------------------------
Mode             : include              Type             : dynamic
Up Time          : 0d 00:00:33          Expires          : 0s
Compat Mode      : MLD Version 2
V1 Host Expires  : 0s                                    :
MVR From-VPLS    : local                MVR To-SAP       : local
MC Standby       : no
-------------------------------------------------------
Source Address
                Up Time      Expires  Type     Fwd/Blk
-------------------------------------------------------
2011::1
                0d 00:00:33  20s      dynamic  Fwd
-------------------------------------------------------------------------------
MLD Group FF04::2
-------------------------------------------------------------------------------
Mode             : include              Type             : dynamic
Up Time          : 0d 00:00:32          Expires          : 0s
Compat Mode      : MLD Version 2
V1 Host Expires  : 0s                                    :
MVR From-VPLS    : local                MVR To-SAP       : local
MC Standby       : no
```

```
-------------------------------------------------
Source Address
                Up Time      Expires  Type     Fwd/Blk
-------------------------------------------------
2011::1
                0d 00:00:33  20s      dynamic  Fwd
-------------------------------------------------------------------------------
Number of groups: 2
===============================================================================
*A:rbae_C#
```

proxy-db

Syntax **proxy-db** [**detail**]
proxy-db group *grp-ipv6-address*

Context show>service>id>mld-snooping

Description This command displays proxy-reporting database entries.

Output The following output is an example of MLD snooping proxy database information.

Sample Output

```
*A:rbae_C# show service id 1 mld-snooping proxy-db
===============================================================================
MLD Snooping Proxy-reporting DB for service 1
===============================================================================
Group Address
                 Mode     Up Time        Num Sources
-------------------------------------------------------------------------------
FF04:db8:1
                 include  0d 00:01:01    1
FF04:db8:2
                 include  0d 00:01:00    1
-------------------------------------------------------------------------------
Number of groups: 2
===============================================================================
*A:rbae_C#


*A:rbae_C# show service id 1 mld-snooping proxy-db detail
===============================================================================
MLD Snooping Proxy-reporting DB for service 1
-------------------------------------------------------------------------------
MLD Group FF04::1
-------------------------------------------------------------------------------
Up Time : 0d 00:01:03                  Mode : include
-------------------------------------------------------------------------
Source Address                                                  Up Time
-------------------------------------------------------------------------
2011:db8:1                                                        0d 00:01:03
-------------------------------------------------------------------------------
MLD Group FF04:db8:2
-------------------------------------------------------------------------------
```

```
Up Time : 0d 00:01:02                   Mode : include
-------------------------------------------------------------------------------
Source Address                                                  Up Time
-------------------------------------------------------------------------------
2011:db8:1                                                        0d 00:01:02
-------------------------------------------------------------------------------
Number of groups: 2
===============================================================================
*A:rbae_C#
```

querier

Syntax **querier**

Context show>service>id>mld-snooping

Description This command displays information about the current querier.

Output The following output is an example of service MLD snooping querier information.

Sample Output

```
*A:rbae_C# show service id 1 mld-snooping querier
===============================================================================
MLD Snooping Querier info for service 1
===============================================================================
Sap Id                  : 2/1/5
IP Address              : FE80:db8:12
Expires                 : 11s
Up Time                 : 0d 00:13:35
Version                 : 2

General Query Interval  : 10s
Query Response Interval : 1.0s
Robust Count            : 2
===============================================================================
*A:rbae_C#
```

static

Syntax **static** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context show>service>id>mld-snooping

Description This command displays MLD snooping static group membership data.

Output The following output is an example of MLD snooping static group membership information.

Sample Output

```
*A:rbae_C# show service id 1 mld-snooping static
===============================================================================
MLD Snooping Static Source Groups for service 1
-------------------------------------------------------------------------------
MLD Snooping Static Source Groups for SDP 36:1
-------------------------------------------------------------------------------
Source
                Group
-------------------------------------------------------------------------------
2011:db8:1
                FF04:db8:2
*
                FF04:db8:3
-------------------------------------------------------------------------------
Static (*,G)/(S,G) entries: 2
===============================================================================
*A:rbae_C#
```

statistics

Syntax **statistics** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context show>service>id>mld-snooping

Description This command displays MLD snooping statistics.

Output The following output is an example of MLD snooping statistics information.

Sample Output

```
*A:rbae_C# show service id 1 mld-snooping statistics
===============================================================================
MLD Snooping Statistics for service 1
===============================================================================
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         109           0             327
Group Queries           0             0             0
Group-Source Queries    0             0             0
V1 Reports              0             0             0
V2 Reports              438           87            0
V1 Done                 0             0             0
Unknown Type            0             N/A           0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length              : 0
Bad MLD Checksum        : 0
Bad Encoding            : 0
No Router Alert         : 0
Zero Source IP          : 0
Wrong Version           : 0
Lcl-Scope Packets       : 0
Rsvd-Scope Packets      : 0
```

```
Send Query Cfg Drops    : 0
Import Policy Drops     : 0
Exceeded Max Num Groups : 0
MCAC Policy Drops       : 0
MCS Failures            : 0

MVR From VPLS Cfg Drops : 0
MVR To SAP Cfg Drops    : 0
===============================================================================
*A:rbae_C#
```

mstp-configuration

Syntax **mstp-configuration**

Context show>service>id

Description This command displays the MSTP specific configuration data. This command is only valid on a management VPLS.

sap

Syntax **sap** *sap-id* [**detail**]

Context show>service>id

Description This command displays information for the SAPs associated with the service.

If no optional parameters are specified, a summary of all associated SAPs is displayed.

Parameters *sap-id* — The ID that displays SAP information.

detail — Displays detailed information for the SAP.

Output The following output is an example of service ID SAP information.

Sample Output

```
*A:ian1# show service id 200 sap 1/1/1:1 detail

===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 200
SAP                : 1/1/1:1                  Encap             : q-tag
Description        : (Not Specified)
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 08/31/2018 16:56:39
Last Mgmt Change   : 08/31/2018 16:56:39
Sub Type           : regular
```

```
Dot1Q Ethertype    : 0x8100                  QinQ Ethertype    : 0x8100
Split Horizon Group: (Not Specified)

Admin MTU          : 1518                    Oper MTU          : 1518
Ingr IP Fltr-Id    : n/a                     Egr IP Fltr-Id    : n/a
Ingr Mac Fltr-Id   : n/a                     Egr Mac Fltr-Id   : n/a
Ingr IPv6 Fltr-Id  : n/a                     Egr IPv6 Fltr-Id  : n/a
qinq-pbit-marking  : both
Egr Agg Rate Limit : max
Q Frame-Based Acct : Disabled                Limit Unused BW   : Disabled
Qinq-vlan-                                   Qinq-vlan-
translation        : None                    translation Ids   : None

Acct. Pol          : None                    Collect Stats     : Disabled

Anti Spoofing      : Ip-Mac                  Dynamic Hosts     : Enabled
Avl Static Hosts   : 1                       Tot Static Hosts  : 1
Calling-Station-Id : n/a

Application Profile: None
Transit Policy     : None
AARP Id            : None

Oper Group         : (none)                  Monitor Oper Grp  : (none)
Host Lockout Plcy  : n/a
Lag Link Map Prof  : (none)
Bandwidth          : Not-Applicable
Oper DCpu Prot Pol*: _default-access-policy

-------------------------------------------------------------------------------
ETH-CFM SAP specifics
-------------------------------------------------------------------------------
Tunnel Faults      : n/a                     AIS               : Disabled
MC Prop-Hold-Timer : n/a
Squelch Levels     : None
Collect Lmm Stats  : Disabled
LMM FC Stats       : None
LMM FC In Prof     : None

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
Ingress qos-policy : 1                       Egress qos-policy : 1
Ingress FP QGrp    : (none)                  Egress Port QGrp  : (none)
Ing FP QGrp Inst   : (none)                  Egr Port QGrp Inst: (none)
Shared Q plcy      : n/a                     Multipoint shared : Disabled
I. Sched Pol       : (Not Specified)
E. Sched Pol       : (Not Specified)
I. Policer Ctl Pol : (Not Specified)
E. Policer Ctl Pol : (Not Specified)
E. HS Sec. Shaper  : (Not Specified)
I. QGrp Redir. List: (Not Specified)
E. QGrp Redir. List: (Not Specified)
-------------------------------------------------------------------------------
Subscriber Management
-------------------------------------------------------------------------------
Admin State        : Up                      MAC DA Hashing    : False
Def Sub-Id         : None
Def Sub-Profile    : basic-sub
```

```
Def SLA-Profile    : basic-sla
Def Inter-Dest-Id  : None
Def App-Profile    : None
Sub-Ident-Policy   : None

Subscriber Limit   : 200
Single-Sub-Parameters (N/A)
 Prof Traffic Only : True
 Non-Sub-Traffic   : N/A

Static host management
MAC learn options  : N/A

-------------------------------------------------------------------------------
Sap Aggregate Stats
-------------------------------------------------------------------------------
                        Packets                 Octets
Ingress
Aggregate Offered     : 0                       0
Aggregate Forwarded   : 0                       0
Aggregate Dropped     : 0                       0

Egress
Aggregate Forwarded   : 0                       0
Aggregate Dropped     : 0                       0
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

                        Packets                 Octets
CPM Ingress           : 0                       0

Forwarding Engine Stats
Dropped               : 0                       0
Received Valid        : 0                       0
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Off. Uncolor          : 0                       0
Off. Managed          : 0                       0

Queueing Stats(Ingress QoS Policy 1)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Queueing Stats(Egress QoS Policy 1)
Dro. In/InplusProf    : 0                       0
Dro. Out/ExcProf      : 0                       0
For. In/InplusProf    : 0                       0
For. Out/ExcProf      : 0                       0
-------------------------------------------------------------------------------
Sap per Queue stats
-------------------------------------------------------------------------------
                        Packets                 Octets

Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                       0
```

```
Off. LowPrio          : 0                      0
Dro. HiPrio           : 0                      0
Dro. LowPrio          : 0                      0
For. InProf           : 0                      0
For. OutProf          : 0                      0

Egress Queue 1
For. In/InplusProf    : 0                      0
For. Out/ExcProf      : 0                      0
Dro. In/InplusProf    : 0                      0
Dro. Out/ExcProf      : 0                      0
===============================================================================
* indicates that the corresponding row element may have been truncated.
*A:ian1#
```

The following table describes show service SAP fields:

Table 27 **Show Service-ID SAP Field Descriptions**

Label	Description
Service Id	The service identifier
SAP	The SAP and qtag
Encap	The encapsulation type of the SAP
Ethertype	The Ethernet type II Ethertype value
Admin State	The administrative state of the SAP
Oper State	The operational state of the SAP
Flags	The conditions that affect the operating status of this SAP
Last Status Change	The time of the most recent operating status change to this SAP
Last Mgmt Change	The time of the most recent management-initiated change to this SAP
Admin MTU	The desired largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented
Oper MTU	The actual largest service frame size (in octets) that can be transmitted through the SAP to the far-end router, without requiring the packet to be fragmented
Ingress qos-policy	The ingress QoS policy ID assigned to the SAP
Egress qos-policy	The egress QoS policy ID assigned to the SAP
Ingress Filter-Id	The ingress filter policy ID assigned to the SAP
Egress Filter-Id	The egress filter policy ID assigned to the SAP

Table 27 **Show Service-ID SAP Field Descriptions (Continued)**

Label	Description (Continued)
Acct. Pol	The accounting policy ID assigned to the SAP
Collect Stats	The collect stats status is enabled or disabled
Dropped	The number of packets and octets dropped due to SAP state, ingress MAC, or IP filter, same segment discard, bad checksum, and so on
Off. HiPrio	The number of high priority packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip.
Off. LowPrio	The number of low priority packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip
Off. Uncolor	The number of uncolored packets and octets, as determined by the SAP ingress QoS policy, offered by the Pchip to the Qchip
Dro. HiPrio	The number of high priority packets and octets, as determined by the SAP ingress QoS policy, dropped by the Qchip due to: MBS exceeded, buffer pool limit exceeded, and so on
Dro. LowPrio	The number of low priority packets and octets, as determined by the SAP ingress QoS policy, dropped by the Qchip due to: MBS exceeded, buffer pool limit exceeded, and so on
Dro. InProf	The number of in-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, and so on
Dro. OutProf	The number of out-of-profile packets and octets discarded by the egress Qchip due to MBS exceeded, buffer pool limit exceeded, and so on
For. InProf	The number of in-profile packets and octets (rate below CIR) forwarded by the egress Qchip
For. OutProf	The number of out-of-profile packets and octets (rate above CIR) forwarded by the egress Qchip

sdp

Syntax **sdp** [*sdp-id* | **far-end** *ip-addr*] [**detail**]

Context show>service>id

Description This command displays information for the SDPs associated with the service.

If no optional parameters are specified, a summary of all associated SDPs is displayed.

Parameters *sdp-id* — Displays the information for the specified SDP ID.

Values 1 to 17407

Default All SDPs

ip-addr — Displays the SDPs matching with the specified far-end IP address.

Default SDPs with any far-end IP address

detail — Displays detailed SDP information.

Output The following output is an example of service SDP information.

Sample Output

```
A:ALA-12#  show service id 9000 sdp 2:22 detail
===============================================================================
Service Destination Point (Sdp Id : 2:22) Details
-------------------------------------------------------------------------------
 Sdp Id 2:22  -(10.10.10.103)
-------------------------------------------------------------------------------
Description     : GRE-10.10.10.103
SDP Id             : 2:22                       Type              : Spoke
Split Horiz Grp : (DSL-group1
VC Type            : Ether                      VC Tag            : n/a
Admin Path MTU     : 4462                       Oper Path MTU     : 4462
Far End            : 10.10.10.103               Delivery          : GRE
Admin State        : Up                         Oper State        : TLDP Down
Ingress Label      : 0                          Egress Label      : 0
Ingress Filter     : n/a                        Egress Filter     : n/a
Last Changed       : 04/11/2007 11:48:20        Signaling         : TLDP

KeepAlive Information :
Admin State        : Disabled                   Oper State        : Disabled
Hello Time         : 10                         Hello Msg Len     : 0
Max Drop Count     : 3                          Hold Down Time    : 10

Statistics            :
I. Fwd. Pkts.      : 0                          I. Dro. Pkts.     : 0
E. Fwd. Pkts.      : 0                          E. Fwd. Octets    : 0

Associated LSP LIST :
SDP Delivery Mechanism is not MPLS
-------------------------------------------------------------------------------
Rstp Service Destination Point specifics
-------------------------------------------------------------------------------
Mac Move           : Disabled
Rstp Admin State   : Up                         Rstp Oper State   : Down
Core Connectivity  : Down
Port Role          : N/A                        Port State        : Discarding
Port Number        : 2049                       Port Priority     : 128
Port Path Cost     : 10                         Auto Edge         : Enabled
Admin Edge         : Disabled                   Oper Edge         : N/A
Link Type          : Pt-pt                      BPDU Encap        : Dot1d
Designated Bridge  : N/A                        Designated Port Id: 0
Active Protocol    : N/A

Fwd Transitions    : 0                          Bad BPDUs rcvd    : 0
```

```
Cfg BPDUs rcvd     : 0                          Cfg BPDUs tx      : 0
TCN BPDUs rcvd     : 0                          TCN BPDUs tx      : 0
RST BPDUs rcvd     : 0                          RST BPDUs tx      : 0
-------------------------------------------------------------------------------
Number of SDPs : 1
===============================================================================
A:ALA-12#
```

The following table describes show service-id SDP output fields:

Table 28 **Show Service-ID SDP Field Descriptions**

Label	Description
Sdp Id	The SDP identifier
Type	The SDP type, spoke or a mesh.
Split Horizon Group	The name of the split horizon group that the SAP belongs to
VC Type	The VC type: ether, vlan, or vpls
VC Tag	The explicit dot1Q value used when encapsulating to the SDP far end
I. Lbl	The VC label used by the far-end device to send packets to this device in this service by the SDP
Admin Path MTU	The operating path MTU of the SDP is equal to the admin path MTU (when one is set) or the dynamically computed tunnel MTU, when no admin path MTU is set (the default case.)
Oper Path MTU	The actual largest service frame size (in octets) that can be transmitted through this SDP to the far-end router, without requiring the packet to be fragmented.
Far End	The IP address of the remote end of the GRE or MPLS tunnel defined by this SDP.
Delivery	The type of delivery used by the SDP: GRE or MPLS
Admin State	The administrative state of this SDP
Oper State	The current status of the SDP
Ingress Label	The label used by the far-end device to send packets to this device in this service by this SDP

interface

Syntax **interface** [{[*ip-address* | *ip-int-name*] [*interface-type*] [**detail**] [**family**]} | **summary**]

Context show>service>id

Description This command displays service interface information.

Parameters *ip-address* — Displays information for the specified IP address.

ip-int-name — Displays information for the specified interface name.

detail — Displays detailed interface information.

family — Specifies to display only information belonging to the address family IPv4 or IPv6. Only these two values are accepted.

summary — Displays detailed interface information.

Output The following output displays an example of service interface information.

Sample Output

```
A:cses-E11>config>service>vprn# show service id 10 interface "gi-2-01" detail
===============================================================================
Interface Table
===============================================================================
-------------------------------------------------------------------------------
Interface
-------------------------------------------------------------------------------
If Name      : gi-2-01
Sub If Name  : si-2
Red If Name  :
Admin State  : Up                              Oper (v4/v6)     : Up/Down
Protocols    : None
-------------------------------------------------------------------------------
Details
-------------------------------------------------------------------------------
Description  : (Not Specified)
If Index     : 6                               Virt. If Index   : 6
Last Oper Chg: 11/27/2012 13:19:28             Global If Index  : 380
Mon Oper Grp : None
Srrp En Rtng : Disabled                        Hold time        : N/A
Group Port   : 1/1/2
TOS Marking  : Trusted                         If Type          : VPRN Grp
SNTP B.Cast  : False
MAC Address  : d2:30:01:01:00:02               Mac Accounting   : Disabled
Ingress stats: Disabled
Arp Timeout  : 14400                           IPv6 Nbr ReachTi*: 30
IP Oper MTU  : 1500                            ICMP Mask Reply  : True
Arp Populate : Disabled                        Host Conn Verify : Disabled
Cflowd       : None
LdpSyncTimer : None
LSR Load Bal*: system
uRPF Chk     : disabled
uRPF Ipv6 Chk: disabled
Rx Pkts      : 0                               Rx Bytes         : 0
Rx V4 Pkts   : 0                               Rx V4 Bytes      : 0
Rx V6 Pkts   : 0                               Rx V6 Bytes      : 0
Tx Pkts      : 32                              Tx Bytes         : 3392
Tx V4 Pkts   : 32                              Tx V4 Bytes      : 3392
Tx V4 Discar*: 0                               Tx V4 Discard By*: 0
Tx V6 Pkts   : 0                               Tx V6 Bytes      : 0
Tx V6 Discar*: 0                               Tx V6 Discard By*: 0
```

```
Proxy ARP Details
Rem Proxy ARP: Disabled                          Local Proxy ARP  : Disabled
Policies     : none

Proxy Neighbor Discovery Details
Local Pxy ND : Disabled
Policies     : none

DHCP no local server

DHCP Details
Description  : (Not Specified)
Admin State  : Down                              Lease Populate   : 0
Gi-Addr      : Not configured                    Gi-Addr as Src Ip: Disabled
Action       : Keep                              Trusted          : Disabled

DHCP Proxy Details
Admin State  : Down
Lease Time   : N/A
Emul. Server : Not configured

Subscriber Authentication Details
Auth Policy  : None

DHCP6 Relay Details
Description  : (Not Specified)
Admin State  : Down                              Lease Populate   : 0
Oper State   : Down                              Nbr Resolution   : Disabled
If-Id Option : None                              Remote Id        : Disabled
Src Addr     : Not configured

DHCP6 Server Details
Admin State  : Down                              Max. Lease States: 8000

ISA Tunnel redundant next-hop information
Static Next-*:
Dynamic Next*:

ICMP Details
Redirects    : Number - 100                      Time (seconds)   - 10
Unreachables : Number - 100                      Time (seconds)   - 10
TTL Expired  : Number - 100                      Time (seconds)   - 10

IPCP Address Extension Details
Peer IP Addr*: Not configured
Peer Pri DNS*: Not configured
Peer Sec DNS*: Not configured
-------------------------------------------------------------------------------
Qos Details
-------------------------------------------------------------------------------

Ing Qos Poli*: (none)                            Egr Qos Policy   : (none)
Ingress FP Q*: (none)                            Egress Port QGrp : (none)
Ing FP QGrp *: (none)                            Egr Port QGrp In*: (none)
-------------------------------------------------------------------------------
Interfaces : 1
===============================================================================
* indicates that the corresponding row element may have been truncated.
```

```
A:cses-E11>config>service>vprn#
```

retailers

Syntax **retailers**

Context show>service>id

Description This command displays service retailer information.

Output The following output displays an example of service retailer information.

Sample Output

```
*A:ALA-48>config# show service id 101 retailers
===============================================================================
Retailers for service 101
===============================================================================
Retailer Svc ID               Num Static Hosts         Num Dynamic Hosts
-------------------------------------------------------------------------------
102                           3                        1
105                           0                        1
-------------------------------------------------------------------------------
Number of retailers : 2
===============================================================================
*A:ALA-48>config#
```

wholesalers

Syntax **wholesalers**

Context show>service>id

Description This command displays service wholesaler information.

Output The following output displays an example of service wholesaler information.

Sample Output

```
*A:ALA-48>config# show service id 102 wholesalers
===============================================================================
Wholesalers for service 102
===============================================================================
Wholesaler Svc ID             Num Static Hosts         Num Dynamic Hosts
-------------------------------------------------------------------------------
101                           3                        1
-------------------------------------------------------------------------------
Number of wholesalers : 1
===============================================================================
*A:ALA-48>config#
```

Wholesaler information can also be displayed in the lease-state context.

```
*A:ALA-48>config# show service id 105 dhcp lease-state  wholesaler 101
===============================================================================
DHCP lease state table, service 105
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LifeTime   Origin   Stdby
-------------------------------------------------------------------------------
Wholesaler 101 Leases
-------------------------------------------------------------------------------
10.3.2.62      00:00:1f:bd:00:c6 lag-1:105           00h00m39s  RADIUS
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
*A:ALA-48>config#
```

split-horizon-group

Syntax **split-horizon-group** [*group-name*]

Context show>service>id

Description This command displays service split horizon groups.

Parameters *group-name* — Specifies a group name up to 32 characters.

Output The following output displays an example of service split horizon group information.

Sample Output

```
A:ALA-1# show service id 700 split-horizon-group
===============================================================================
Service: Split Horizon Group
===============================================================================
Name                            Description
-------------------------------------------------------------------------------
R DSL-group1                    Split horizon group for DSL
-------------------------------------------------------------------------------
R = Residential Split Horizon Group
No. of Split Horizon Groups: 1
===============================================================================
A:ALA-1#

A:ALA-1# show service id 700 split-horizon-group DSL-group1
===============================================================================
Service: Split Horizon Group
===============================================================================
Name                            Description
-------------------------------------------------------------------------------
R DSL-group1                    Split horizon group for DSL
-------------------------------------------------------------------------------
Associations
```

```
-------------------------------------------------------------------------------
SAP                              1/1/3:1

SDP                              108:1
SDP                              109:1
-------------------------------------------------------------------------------
R = Residential Split Horizon Group
SAPs Associated : 1              SDPs Associated : 2
===============================================================================
A:ALA-1#
```

static-host

Syntax **static-host** [**sap** *sap-id*] [**wholesaler** *service-id*] [**port** *port-id*][**inter-dest-id** *intermediate-destination-id*] [**detail**] [**mac** *ieee-address*] [**ip-address** *ip-prefix*[*/prefix-length*]]
static-host [**sap** *sap-id*] [**wholesaler** *service-id*] [**port** *port-id*] **no-inter-dest-id** [**detail**]
static-host summary

Context show>service>id

Description This command displays static hosts configured on this service.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

intermediate-destination-id — Specifies the intermediate destination identifier which is encoded in the identification strings up to 32 characters.

summary — Displays summary static host information.

detail — Displays detailed static host information.

service-id — The service ID of the wholesaler (applies only to the 7750 SR).

Values 1 to 2147483647

ieee-address — Specifies the MAC address.

*ip-prefix***[***/prefix-length***]** — Specifies the IP address.

Output The following output displays an example of service static host information.

Sample Output

```
*A:ALA-48# show service id 88 static-host
===============================================================================
Static Hosts for service 88
===============================================================================
Sap                          IP Address      Configured MAC    Dynamic MAC
  Subscriber                                   Admin State       Fwding State
-------------------------------------------------------------------------------
1/2/20:0                     10.10.10.104    N/A               N/A
   N/A                                         Down              Not Fwding
3/2/4:50/5                   10.144.145.1    N/A               N/A
   N/A                                         Up                Fwding
-------------------------------------------------------------------------------
```

```
Number of static hosts : 2
===============================================================================
*A:ALA-48#
```

stp

Syntax **stp** [**detail**]

Context show>service>id

Description Displays information for the spanning tree protocol instance for the service.

Parameters **detail** — Displays detailed information.

Output The following output displays an example of service STP information.

Sample Output

```
A:ALA-12# show service id 1 stp
===============================================================================
Spanning Tree Information
===============================================================================
VPLS Spanning Tree Information
-------------------------------------------------------------------------------
RSTP Admin State  : Up                  RSTP Oper State   : Down
Core Connectivity : Down
Bridge-id         : 04:67:ff:00:00:01

Hold Timer        : 1                   Bridge fwd delay  : 15
Bridge Hello time : 1                   Bridge max age    : 20
Bridge priority   : 1                   Topology change   : Inactive
Last Top. change  : 0d 00:00:00         Top. change count : 0

Root bridge-id    : 00:03:fa:00:00:00

Root path cost    : 1                   Root forward delay: 15
Root hello time   : 1                   Root max age      : 20
Root priority     : 0                   Root port         : vcp

-------------------------------------------------------------------------------
Spanning Tree Specifics
-------------------------------------------------------------------------------
SAP Identifier    : 1/1/7:0             RSTP State        : Down
STP Port State    : Forwarding          BPDU encap        : dot1d
Port Number       : 2048                Priority          : 128
Cost              : 10                  Fast Start        : Disabled
Designated Port   : 34816               Designated Bridge : 02:fa:00:04:54:01
===============================================================================
A:ALA-12#
```

The following is an example for the 7450 ESS:

```
ALA-1# show service id 1 stp
===============================================================================
```

```
Inherited Rstp State (from mVPLS), Service 1
===============================================================================
Sap/Spoke Id       Oper-     Prune-  Port-       Managed by  Managed by
                   State     State   State       Service     Sap/spoke
-------------------------------------------------------------------------------
1/1/1:2            Down      Pruned  Discarding  101         1/1/1:1
100:1              Up        Pruned  Discarding  100         100:100
===============================================================================
```

The following table describes show service-id STP output fields:

Table 29 **Show Service-ID STP Output Field Descriptions**

Label	Description
RSTP Admin State	The administrative state of the Rapid Spanning Tree Protocol instance associated with this service
Core Connectivity	The connectivity status to the core
RSTP Oper State	The operational state of the Rapid Spanning Tree Protocol instance associated with this service. This field is applicable only when STP is enabled on the router.
Bridge-id	The MAC address used to identify this bridge in the network.
Hold Time	The interval length during which no more than two Configuration BPDUs shall be transmitted by this bridge
Bridge fwd delay	The speed of bridge state changes when moving toward the forwarding state
Bridge Hello time	The amount of time between the transmission of Configuration BPDUs
Bridge max age	The maximum age of spanning tree protocol information learned from the network on any port before it is discarded. This is the actual value that this bridge is currently using.
Bridge priority	The priority of the spanning tree protocol instance associated with this service
Topology change	Topology change is currently in progress.
Last Top. change	The time (in hundredths of a second) since the last time a topology change was detected by the Spanning Tree Protocol instance associated with this service
Top. change count	The total number of topology changes detected by the Spanning Tree Protocol instance associated with this service since the management entity was last reset or initialized
Root bridge-id	The Root bridge-id

Table 29 **Show Service-ID STP Output Field Descriptions (Continued)**

Label	Description (Continued)
Root path cost	The Root path cost
Root forward delay	The Root forward delay
Root hello time	The Root hello time
Root max age	The Root maximum age
Root priority	The Root priority level
Root port	The Root port number
SAP Identifier	The SAP identifier
RSTP State	The RSTP state
STP Port State	The STP port state
BPDU encap	The BPDU encapsulation
Port Number	The port number
Priority	The priority level
Cost	The cost
Fast Start	The fast start state
Designated Port	The designated port
Designated Bridge	The designated bridge
Sap/Spoke Id	The SAP or spoke ID
Oper-State	The operational state
Prune-State	The prune state
Port-State	The port state
Managed by Service	The managed service
Managed by Sap/ spoke	The managed SAP or spoke

subscriber-hosts

Syntax **subscriber-hosts** [**sap** *sap-id*] [**ip** *ip-prefix/prefix-length*] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**app-profile** *app-profile-name*] [**wholesaler** *service-id*] [**address-origin** *address-origin*] [**detail**] [**statistics**] [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]

Context show>service>id

Description This command displays subscriber host information.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

ip-prefix/prefix-length — Shows information for the specified IP address and mask.

Values *ipv4-prefix* — a.b.c.d (host bits must be 0)
ipv4-prefix-length — 0 to 32
ipv6-prefix — x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x — [0..FFFF]H
d — [0..255]D
ipv6-prefix-length — 0 to 128

ieee-address — Displays information only for the specified 48-bit MAC address. The MAC address can be expressed in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.

sub-profile-name — Displays an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config**>**subscr-mgmt**>**sub-profile** context.

sla-profile-name — Displays this optional parameter to specify an existing SLA profile name, up to 32 characters.

app-profile-name — Displays this optional parameter to specify an existing SLA app profile name, up to 32 characters.

service-id — The ID that uniquely identifies a service.

Values service-id: 1 to 214748364
svc-name: A string up to 64 characters.

address-origin — displays this optional parameter to specify an address origin.

Values aaa, dynamic, static, bonding

detail — Displays detailed information.

statistics — Displays the HTTP redirect statistics for the host. The redirect statistics is divided into HTTP GET and HTTP response for redirection.

spi-sharing-type — Displays subscriber host information for sessions with the specified SPI sharing type.

Values per-group, per-session-ipoe, per-session-ppp

spi-sharing-id — Displays subscriber host information for sessions with the specified SPI sharing type and sharing identifier. The SPI group ID is:

- a group ID when sharing type is per-group

- an internal PPP or IPoE sharing ID when the sharing type is per-PPP or per-IPoE session

Values 0 to 4294967295

Output The following output display is an example of subscriber host information.

Sample Output

```
*A:Dut-C>#  show service id 3 subscriber-hosts
===============================================================
Subscriber Host table
===============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID     Origin       Fwding State
---------------------------------------------------------------
2/1/5:2                    TEACAHEH74
  10.11.1.61
    00:80:00:00:00:0a          N/A               ARP-Host        Fwding
[pw-11:11]                  VIACAHEH74
  10.11.1.2
    00:00:11:11:01:02          N/A               ARP-Host        Fwding
[pw-11:12]                  pw-11:12
  10.11.1.3
    00:00:11:11:01:03          N/A               ARP-Host        Fwding
[pw-11:13]                  pw-11:13
  10.11.1.4
    00:00:11:11:01:04          N/A               ARP-Host        Fwding
[pw-22:22]                  XMACAHEH74
  10.22.1.2
    00:00:22:22:01:02          N/A               ARP-Host        Fwding

[pw-33:33]                  IUASAHEH74
  10.33.1.2
    00:00:33:33:01:02          N/A               ARP-Host        Fwding
---------------------------------------------------------------
Number of subscriber hosts: 6
===============================================================
*A:Dut-C>#


A:Dut-A# show service id 100 subscriber-hosts ip 10.100.1.5
===============================================================================
Subscriber Host table
===============================================================================
Sap                    IP Address      MAC Address       Origin(*) Subscriber
-------------------------------------------------------------------------------
1/2/1:102              10.100.1.5      00:10:00:00:00:03 -/D/- nokia_100
-------------------------------------------------------------------------------
Number of subscriber hosts : 1
===============================================================================
(*) S=Static Host, D=DHCP Lease, N=Non-Sub-Traffic
A:Dut-A#
```

The following table describes show service-id subscriber hosts output fields:

Table 30 **Show Service-ID Subscriber Hosts Field Descriptions**

Label	Description
Sap	The SAP ID number
IP Address	The IP address
MAC Address	The MAC address
PPPoE-SID	The PPPoE-SID of the originating subscriber
Origin Subscriber	The ID of the originating subscriber
Fwding State	The forwarding state
Number of subscriber hosts	The number of subscriber hosts

sdp

Syntax **sdp** *sdp-id* **pw-port** [*pw-port-id*]
sdp *sdp-id* **pw-port**
sdp *sdp-id* **pw-port** [*pw-port-id*] [**statistics**]
sdp [**consistent** | **inconsistent** | **na**] **egressifs**
sdp *sdp-id* **keep-alive-history**
sdp far-end *ip-address* | *ipv6-address* **keep-alive-history**
sdp [***sdp-id***] **detail**
sdp far-end *ip-address* | *ipv6-address* **detail**

Context show>service>sdp

Description This command displays SDP information.

If no optional parameters are specified, a summary SDP output for all SDPs is displayed.

Parameters *sdp-id* — The SDP ID for which to display information.

Values 1 to 17407

Default All SDPs

pw-port-id — Displays the SAP identifier for PW-SAPs

Values 1 to 10239

ip-address — Displays only SDPs matching with the specified far-end IP address

Default SDPs with any far-end IP address

detail — Displays detailed SDP information

Default SDP summary output

keep-alive-history — Displays the last fifty SDP keepalive events for the SDP

Default SDP summary output

Output The following output displays service SDP information.

Sample Output

```
===============================================================================
*A:ALA-12>config>service# show service sdp 1 pw-port
===============================================================================
Service Destination Point (sdp Id 1 Pw-Port)
===============================================================================
Pw-port   VC-Id     Adm     Encap       Opr     VC Type     Egr       Monitor
                                                           Shaper    Oper
                                                           VPort     Group
-------------------------------------------------------------------------------
1         1         up      dot1q       up      ether
2         2         up      qinq        up      ether
3         3         up      dot1q       up      ether
4         4         up      qinq        up      ether
-------------------------------------------------------------------------------
Entries found : 4
===============================================================================

*A:ALA-12>config>service# show service sdp 1 pw-port 3
===============================================================================
Service Destination Point (Sdp Id 1 Pw-Port 3)
===============================================================================
SDP Binding port     : lag-1
VC-Id                : 3                    Admin Status       : up
Encap                : dot1q                Oper Status        : up
VC Type              : ether
Oper Flags           : (Not Specified)
Monitor Oper-Group   : (Not Specified)
===============================================================================

*A:ALA-12>config>service# show service sdp 1 pw-port 3 statistics
===============================================================================
Service Destination Point (Sdp Id 1 Pw-Port 3)
===============================================================================
SDP Binding port     : lag-1
VC-Id                : 3                    Admin Status       : up
Encap                : dot1q                Oper Status        : up
VC Type              : ether
Oper Flags           : (Not Specified)
Monitor Oper-Group   : (Not Specified)

Statistics            :
I. Fwd. Pkts.        : 0                    I. Dro. Pkts.      : 0
I. Fwd. Octs.        : 0                    I. Dro. Octs.      : 0
E. Fwd. Pkts.        : 0                    E. Fwd. Octets     : 0
===============================================================================
```

subscriber-using

Syntax **subscriber-using** [**service-id** *service-id*] [**sap-id** *sap-id*] [**interface** *ip-int-name*] [**ip** *ip-address*[/*mask*]] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**app-profile** *app-profile-name*] [**port** *port-id*] [**no-inter-dest-id** | **inter-dest-id** *intermediate-destination-id*]

Context show>service

Description This command displays selective subscriber information using specific options.

Parameters *service-id* — Displays information for the specifies ID that uniquely identifies a service.

sap-id — Displays the physical port identifier portion of the SAP definition.

ip-int-name — Shows DHCP statistics on the specified interface.

port-id — Indicates the SAP or SDP for which this entry contains information.

ip-address[/*mask*] — Shows information for the specified IP address and mask.

ieee-address — Displays information only for the specified 48-bit MAC address. The MAC address can be expressed in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.

sub-profile-name — Displays an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.*sla-profile-name*.

sla-profile — Displays this optional parameter to specify an existing SLA profile name.

app-profile — Displays the application specified profile.

intermediate-destination-id — Indicates the intermediate destination identifier received from either the DHCP or the RADIUS server or the local user database.

Output The following output displays an example of selective subscriber information using specific options.

Sample Output

```
A:Dut-A# show service subscriber-using service-id 100
===============================================================
Subscribers
===============================================================
Subscriber                      Sub Profile
---------------------------------------------------------------
nokia_100                     sub_prof100
---------------------------------------------------------------
Matching Subscribers : 1
===============================================================
A:Dut-A#

A:Dut-A# show service subscriber-using
===============================================================
Subscribers
===============================================================
```

```
Subscriber                        Sub Profile
-------------------------------------------------------------------
nokia_100                      sub_prof100
nokia_110                      sub_prof110
nokia_120                      sub_prof120
nokia_130                      sub_prof130
nokia_80                       sub_prof80
nokia_90                       sub_prof90
client_PC1                       sub_profPC1
static                           sub_default
-------------------------------------------------------------------
Matching Subscribers : 8
===================================================================
A:Dut-A#
```

redundancy

Syntax **redundancy**

Context show

Description This command enables the context to show multi-chassis redundancy information.

multi-chassis

Syntax **multi-chassis all**
multi-chassis mc-lag peer *ip-address* [**lag** *lag-id*]
multi-chassis mc-lag [**peer** *ip-address* [**lag** *lag-id*]] **statistics**
multi-chassis sync [**peer** *ip-address*] [**detail**]
multi-chassis sync [**peer** *ip-address*] **statistics**

Context show>redundancy

Description This command displays multi-chassis redundancy information.

Parameters **all** — Displays all multi-chassis information.

mc-lag — Displays multi-chassis LAG information.

ip-address — Displays the address of the multi-chassis peer.

lag-id — Displays the specified LAG ID on this system that forms an multi-chassis LAG configuration with the indicated peer.

statistics — Displays statistics for the multi-chassis peer.

sync — Displays synchronization information.

detail — Displays detailed information.

Output The following output displays an example of multi-chassis redundancy information.

Sample Output

```
A:pc1# show redundancy multi-chassis all
===============================================================================
Multi-Chassis Peers
===============================================================================
Peer IP          Src IP          Auth            Peer Admin
  MCS Admin       MCS Oper        MCS State       MC-LAG Admin    MC-LAG Oper
-------------------------------------------------------------------------------
10.10.10.102     10.10.10.101    hash            Enabled
  Enabled         Enabled         inSync          Enabled         Enabled
10.10.20.1       0.0.0.0         None            Disabled
  --              --              --              Disabled        Disabled
===============================================================================
A:pc1#


*A:Dut-C# show redundancy multi-chassis mc-lag peer 10.10.10.1
===============================================================================
Multi-Chassis MC-Lag Peer 10.10.10.1
===============================================================================
Last State chg: 09/24/2007 07:58:03
Admin State: Up      Oper State   : Up
KeepAlive: 10 deci-seconds     Hold On Ngbr Failure : 3
-------------------------------------------------------------------------------
Lag Id Lacp Key Remote Lag Id System Id  Sys Prio Last State Changed
-------------------------------------------------------------------------------
1      326661      00:00:00:33:33:33  32888  09/24/2007 07:56:35
-------------------------------------------------------------------------------
Number of LAGs : 1
===============================================================================
*A:Dut-C#

A:pc1# show redundancy multi-chassis mc-lag statistics
===============================================================================
Multi-Chassis Statistics
===============================================================================
Packets Rx                        : 129816
Packets Rx Keepalive              : 129798
Packets Rx Config                 : 3
Packets Rx Peer Config            : 5
Packets Rx State                  : 10
Packets Dropped KeepaliveTask     : 0
Packets Dropped Packet Too Short  : 0
Packets Dropped Verify Failed     : 0
Packets Dropped Tlv Invalid Size  : 0
Packets Dropped Out of Seq        : 0
Packets Dropped Unknown Tlv       : 0
Packets Dropped Tlv Invalid LagId : 0
Packets Dropped MD5               : 0
Packets Dropped Unknown Peer      : 0
Packets Tx                        : 77918
Packets Tx Keepalive              : 77879
Packets Tx Config                 : 6
Packets Tx Peer Config            : 26
Packets Tx State                  : 7
Packets Tx Failed                 : 0
===============================================================================
A:pc1#
```

```
A:pc1# show redundancy multi-chassis mc-lag peer 10.10.10.102 lag 2 statistics
===============================================================================
Multi-Chassis Statistics, Peer 10.10.10.102 Lag 2
===============================================================================
Packets Rx Config                : 1
Packets Rx State                 : 4
Packets Tx Config                : 2
Packets Tx State                 : 3
Packets Tx Failed                : 0
===============================================================================
A:pc1#


A:pc1#show redundancy multi-chassis mc-lag peer 10.10.10.102 statistics
===============================================================================
Multi-Chassis Statistics, Peer 10.10.10.102
===============================================================================
Packets Rx                       : 129918
Packets Rx Keepalive             : 129900
Packets Rx Config                : 3
Packets Rx Peer Config           : 5
Packets Rx State                 : 10
Packets Dropped State Disabled   : 0
Packets Dropped Packets Too Short : 0
Packets Dropped Tlv Invalid Size  : 0
Packets Dropped Tlv Invalid LagId : 0
Packets Dropped Out of Seq       : 0
Packets Dropped Unknown Tlv      : 0
Packets Dropped MD5              : 0
Packets Tx                       : 77979
Packets Tx Keepalive             : 77940
Packets Tx Peer Config           : 26
Packets Tx Failed                : 0
===============================================================================
A:pc1#


A:pc1# show redundancy multi-chassis sync
===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 10.10.10.102
Description             : CO1
Authentication          : Enabled
Source IP Address       : 10.10.10.101
Admin State             : Enabled
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     :
Sync Admin State        : Up
Sync Oper State         : Up
DB Sync State           : inSync
Num Entries             : 0
Lcl Deleted Entries     : 0
```

```
Alarm Entries           : 0
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 10.10.20.1
Authentication          : Disabled
Source IP Address       : 0.0.0.0
Admin State             : Disabled
===============================================================================
A:pc1#

pc1# show redundancy multi-chassis sync peer 10.10.10.102
===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 10.10.10.102
Description             : CO1
Authentication          : Enabled
Source IP Address       : 10.10.10.101
Admin State             : Enabled
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     :
Sync Admin State        : Up
Sync Oper State         : Up
DB Sync State           : inSync
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
===============================================================================
MCS Application Stats
===============================================================================
Application             : igmp
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
-------------------------------------------------------------------------------
Application             : igmpSnooping
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
-------------------------------------------------------------------------------
```

```
Application             : subMgmt
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
-------------------------------------------------------------------------------
Application             : srrp
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
===============================================================================
A:pc1#

A:pc1# show redundancy multi-chassis sync peer 10.10.10.102 detail
===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 10.10.10.102
Description             : CO1
Authentication          : Enabled
Source IP Address       : 10.10.10.101
Admin State             : Enabled
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     :
Sync Admin State        : Up
Sync Oper State         : Up
DB Sync State           : inSync
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
===============================================================================
MCS Application Stats
===============================================================================
Application             : igmp
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
-------------------------------------------------------------------------------
Application             : igmpSnooping
Num Entries             : 0
Lcl Deleted Entries     : 0
```

```
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
-------------------------------------------------------------------------------
Application             : subMgmt
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
-------------------------------------------------------------------------------
Application             : srrp
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
===============================================================================
Ports synced on peer 10.10.10.102
===============================================================================
Port/Encap                    Tag
-------------------------------------------------------------------------------
1/1/1
  1-2                         r1
===============================================================================
A:pc1#

A:pc1# show redundancy multi-chassis sync statistics
===============================================================================
Multi-chassis Peer Sync Stats
===============================================================================
Peer IP Address         : 10.10.10.102
Packets Tx Total        : 511
Packets Tx Hello        : 510
Packets Tx Data         : 0
Packets Tx Other        : 1
Packets Tx Error        : 0
Packets Rx Total        : 511
Packets Rx Hello        : 510
Packets Rx Data         : 0
Packets Rx Other        : 1
Packets Rx Error        : 0
Packets Rx Header Err   : 0
Packets Rx Body Err     : 0
Packets Rx Seq Num Err  : 0
===============================================================================
Peer IP Address         : 10.10.20.1
Packets Tx Total        : 0
Packets Tx Hello        : 0
Packets Tx Data         : 0
Packets Tx Other        : 0
Packets Tx Error        : 0
Packets Rx Total        : 0
```

```
Packets Rx Hello        : 0
Packets Rx Data         : 0
Packets Rx Other        : 0
Packets Rx Error        : 0
Packets Rx Header Err   : 0
Packets Rx Body Err     : 0
Packets Rx Seq Num Err  : 0
===============================================================================
A:pc1#

A:pc1# show redundancy multi-chassis sync peer 10.10.10.102 statistics
===============================================================================
Multi-chassis Peer Sync Stats
===============================================================================
Peer IP Address         : 10.10.10.102
Packets Tx Total        : 554
Packets Tx Hello        : 553
Packets Tx Data         : 0
Packets Tx Other        : 1
Packets Tx Error        : 0
Packets Rx Total        : 554
Packets Rx Hello        : 553
Packets Rx Data         : 0
Packets Rx Other        : 1
Packets Rx Error        : 0
Packets Rx Header Err   : 0
Packets Rx Body Err     : 0
Packets Rx Seq Num Err  : 0
===============================================================================
A:pc1#
```

mc-ipsec

Syntax **mc-ipsec peer** *addr* [**tunnel-group** *group-id*]

Context show>redundancy>multi-chassis

Description This command displays the IPsec multi-chassis states. Optionally, only the states of the specified tunnel-groups are displayed.

Parameters *addr* — Specifies the address of the peer.

group-id — Specifies the tunnel-group ID.

Output The following output displays an example of redundant multi-chassis peer information.

Sample Output

```
===============================================================================
Multi-Chassis MC-IPSec
===============================================================================
Peer Name       : (Not Specified)
Peer Addr       : 2.2.2.2
Keep Alive Intvl: 10                  Hold on Nbr Fail     : 3
BFD Intf Name   : None
```

```
BFD Dest Addr   :
Last update     : 03/20/2012 22:48:55
===============================================================================
Multi-Chassis IPsec Multi Active Tunnel-Group Table
===============================================================================
ID             Peer Group     Priority  Preempt   Admin State    Mastership
-------------------------------------------------------------------------------
1              1              100       Disabled  Up             eligible
-------------------------------------------------------------------------------
Multi Active Tunnel Group Entries found: 1
===============================================================================
```

mc-ring

Syntax **mc-ring peer** *ip-address* **statistics**
mc-ring peer *ip-address* [**ring** *sync-tag* [**detail** | **statistics**]]
mc-ring peer *ip-address* **ring** *sync-tag* **ring-node** [*ring-node-name* [**detail** | **statistics**]]
mc-ring global-statistics

Context show>redundancy>multi-chassis

Description This command displays multi-chassis ring information.

Parameters *ip-address* — Specifies the address of the multi-chassis peer to display.

sync-tag — Specifies a synchronization tag to be displayed that was used while synchronizing this port with the multi-chassis peer.

ring-node-name — Specifies a ring-node name.

global-statistics — Displays global statistics for the multi-chassis ring.

detail — Displays detailed peer information for the multi-chassis ring.

Output The following output is an example of mc-ring peer information.

Sample Output

```
show redundancy multi-chassis mc-ring peer 10.0.0.2 ring ring11 detail
===============================================================================
Multi-Chassis MC-Ring Detailed Information
===============================================================================
Peer           : 10.0.0.2
Sync Tag       : ring11
Port ID        : 1/1/3
Admin State    : inService
Oper State     : connected
Admin Change   : 01/07/2008 21:40:07
Oper Change    : 01/07/2008 21:40:24
Failure Reason : None
-------------------------------------------------------------------------------
In Band Control Path
-------------------------------------------------------------------------------
Service ID     : 10
```

```
Interface Name : to_an1
Oper State     : connected
Dest IP        : 10.10.0.2
Src  IP        : 10.10.0.1
-------------------------------------------------------------------------------
VLAN Map B Path Provisioned
-------------------------------------------------------------------------------
range 13-13
range 17-17
-------------------------------------------------------------------------------
VLAN Map Excluded Path Provisioned
-------------------------------------------------------------------------------
range 18-18
-------------------------------------------------------------------------------
VLAN Map B Path Operational
-------------------------------------------------------------------------------
range 13-13
range 17-17
-------------------------------------------------------------------------------
VLAN Map Excluded Path Operational
-------------------------------------------------------------------------------
range 18-18
===============================================================================


*A:ALA-48>show>redundancy>multi-chassis# mc-ring peer 192.251.10.104
===============================================================================
MC Ring entries
===============================================================================
Sync Tag                        Oper State      Failure Reason
-------------------------------------------------------------------------------
No. of MC Ring entries: 0
===============================================================================

show redundancy multi-chassis mc-ring peer 10.0.0.2
===============================================================================
MC Ring entries
===============================================================================
Sync Tag                        Oper State      Failure Reason
-------------------------------------------------------------------------------
ring11                          connected       None
ring12                          shutdown        None
-------------------------------------------------------------------------------
No. of MC Ring entries: 4
===============================================================================


show redundancy multi-chassis mc-ring peer 10.0.0.2 ring ring11 ring-node an1 detail
===============================================================================
Multi-Chassis MC-Ring Node Detailed Information
===============================================================================
Peer           : 10.0.0.2
Sync Tag       : ring11
Node Name      : an1
Oper State Loc : connected
Oper State Rem : notTested
In Use         : True
Admin Change   : 01/07/2008 21:40:07
Oper Change    : 01/07/2008 21:40:25
```

```
Failure Reason : None
-------------------------------------------------------------------------------
Ring Node Connectivity Verification
-------------------------------------------------------------------------------
Admin State    : inService
Service ID     : 11
VLAN Tag       : 11
Dest IP        : 10.11.3.1
Src  IP        : None
Interval       : 1 minutes
Src MAC        : None
===============================================================================

show redundancy multi-chassis mc-ring peer 10.0.0.2 ring ring11 ring-node
===============================================================================
MC Ring Node entries
===============================================================================
Name                            Loc Oper St.      Failure Reason
  In Use                          Rem Oper St.
-------------------------------------------------------------------------------
an1                             connected         None
  Yes                             notTested
an2                             connected         None
  Yes                             notTested
-------------------------------------------------------------------------------
No. of MC Ring Node entries: 2
===============================================================================
```

The following table describes multi-chassis ring peer output fields

Table 31 **Show mc-ring peer ip-address ring Output Field Descriptions**

Label	Description
Sync Tag	The synchronization tag that was used while synchronizing this port with the multi-chassis peer

Table 31 **Show mc-ring peer ip-address ring Output Field Descriptions**

Label	Description (Continued)
Oper State	noPeer — The peer has no corresponding ring configured. connected — The inband control connection with the peer is operational. broken — The inband control connection with the peer has timed out. conflict — The inband control connection with the peer has timed out but the physical connection is still OK; the failure of the inband signaling connection is caused by a misconfiguration. For example, a conflict between the configuration of this system and its peer, or a misconfiguration on one of the ring access node systems. testingRing — The inband control connection with the peer is being set up. Waiting for result. waitingForPeer — Verifying if this ring is configured on the peer. configErr — The ring is administratively up, but a configuration error prevents it from operating properly. halfBroken — The inband control connection indicates that the ring is broken in one direction (towards the peer). localBroken — The inband control connection with the peer is known to be broken due to local failure or local administrative action. shutdown — The ring is shutdown.
Failure Reason	The failure reason.
No. of MC Ring entries	The number of MC ring entries.

```
*A:ALA-48>show>redundancy>multi-chassis# mc-ring peer 192.251.10.104 statistics
===============================================================================
MC Ring statistics for peer 192.251.10.104
===============================================================================
Message                                                 Received     Transmitted
-------------------------------------------------------------------------------
MCS ID Request                                          0            0
MCS ID Response                                         0            0
Ring Exists Request                                     0            0
Ring Exists Response                                    0            0
Keepalive                                               0            0
-------------------------------------------------------------------------------
Total                                                   0            0
===============================================================================
*A:ALA-48>show>redundancy>multi-chassis#
```

The following tables describe multi-chassis ring peer output fields

Table 32 **Multi-chassis Ring Peer Output Field Descriptions**

Label	Description
Message	The message type
Received	The number of valid MC-Ring signaling messages received from the peer
Transmitted	The number of valid MC-Ring signaling messages transmitted from the peer
MCS ID Request	The number of valid MCS ID requests were received from the peer
MCS ID Response	The number of valid MCS ID responses were received from the peer
Ring Exists Request	The number of valid 'ring exists' requests were received from the peer
Ring Exists Response	The number of valid ring exists' responses were received from the peer
Keepalive	The number of valid MC-Ring control packets of type 'keepalive' were received from the peer

The following table describes multi-chassis ring node fields.

Table 33 **Show mc-ring ring-node Output Field Descriptions**

Label	Description
Oper State	The state of the connection verification (both local and remote) notProvisioned — Connection verification is not provisioned. configErr — Connection verification is provisioned but a configuration error prevents it from operating properly. notTested — Connection verification is administratively disabled or is not possible in the current situation. testing — Connection Verification is active, but no results are yet available. connected — The ring node is reachable. Oper State (Cont) disconnected — Connection verification has timed out.
In Use	“True” if the ring node is referenced on an Epipe or as an inter-dest-id on a static host or dynamic lease

show mc-ring global-statistics Output

The following table describes mc-ring global statistics information.

Sample Output

```
*A:ALA-48>show>redundancy>multi-chassis# mc-ring global-statistics
===============================================================================
Global MC Ring statistics
===============================================================================
Rx                          : 0
Rx Too Short                : 0
Rx Wrong Authentication     : 0
Rx Invalid TLV              : 0
Rx Incomplete               : 0
Rx Unknown Type             : 0
Rx Unknown Peer             : 0
Rx Unknown Ring             : 0
Rx Unknown Ring Node        : 0
Tx                          : 36763
Tx No Buffer                : 0
Tx Transmission Failed      : 0
Tx Unknown Destination      : 0
Missed Configuration Events : 0
Missed BFD Events           : 0
===============================================================================
*A:ALA-48>show>redundancy>multi-chassis#
```

The following table describes mc-ring global statistics fields.

Table 34 **Show mc-ring global-statistics Output Field Descriptions**

Label	Description
Rx	The number of MC-ring signaling packets were received by this system
Rx Too Short	The number of MC-ring signaling packets were received by this system that were too short
Rx Wrong Authentication	The number of MC-ring signaling packets were received by this system with invalid authentication
Rx Invalid TLV	The number of MC-ring signaling packets were received by this system with invalid TLV
Rx Incomplete	The number of MC-ring signaling packets were received by this system that were incomplete
Rx Unknown Type	The number of MC-ring signaling packets were received by this system that were of unknown type
Rx Unknown Peer	The number of MC-ring signaling packets were received by this system that were related to an unknown peer

Table 34 **Show mc-ring global-statistics Output Field Descriptions**

Label	Description (Continued)
Rx Unknown Ring	The number of MC-ring signaling packets were received by this system that were related to an unknown ring
Rx Unknown Ring Node	The number of MC-ring signaling packets were received by this system that were related to an unknown ring node
Tx	The number of MC-ring signaling packets were transmitted by this system
Tx No Buffer	The number of MC-ring signaling packets could not be transmitted by this system due to a lack of packet buffers
Tx Transmission Failed	The number of MC-ring signaling packets could not be transmitted by this system due to a transmission failure
Tx Unknown Destination	The number of MC-ring **unknown destination** signaling packets were transmitted by this system
Missed Configuration Events	The number of missed configuration events on this system
Missed BFD Events	The number of missed BFD events on this system

2.11.2.3.1 IGMP Snooping Show Commands

igmp-snooping

Syntax **igmp-snooping**

Context show>service>id

Description This command enables the context to display IGMP snooping information.

all

Syntax **all**

Context show>service>id>igmp-snooping

Description Displays detailed information for all aspects of IGMP snooping on the VPLS service.

Output The following output displays an example of IGMP snooping information.

Sample Output

The following example applies to the 7750 SR:

```
*A:ALA-48>show>service>id>igmp-snooping>snooping# all
===============================================================================
IGMP Snooping info for service 750
===============================================================================
IGMP Snooping Base info
-------------------------------------------------------------------------------
Admin State : Up
Querier     : No querier found
-------------------------------------------------------------------------------
Sap/Sdp                Oper      MRtr  Send      Max Num   Num
Id                     State     Port  Queries   Groups    Groups
-------------------------------------------------------------------------------
sap:1/1/7:0            Down      No    Disabled  No Limit  0
sdp:1:22               Down      No    Disabled  No Limit  0
sdp:8:750              Down      No    Disabled  No Limit  0
-------------------------------------------------------------------------------
IGMP Snooping Querier info
-------------------------------------------------------------------------------
No querier found for this service.
-------------------------------------------------------------------------------
IGMP Snooping Multicast Routers
-------------------------------------------------------------------------------
MRouter          Sap/Sdp Id             Up Time        Expires   Version
-------------------------------------------------------------------------------
Number of mrouters: 0
-------------------------------------------------------------------------------
IGMP Snooping Proxy-reporting DB
-------------------------------------------------------------------------------
Group Address    Mode     Type     Up Time        Expires  Num Src
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
IGMP Snooping SAP 1/1/7:0 Port-DB
-------------------------------------------------------------------------------
Group Address    Mode     Type     Up Time        Expires  Num Src
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
IGMP Snooping SDP 1:22 Port-DB
-------------------------------------------------------------------------------
Group Address    Mode     Type     Up Time        Expires  Num Src
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
IGMP Snooping SDP 8:750 Port-DB
-------------------------------------------------------------------------------
Group Address    Mode     Type     Up Time        Expires  Num Src
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
IGMP Snooping Static Source Groups
-------------------------------------------------------------------------------
IGMP Snooping Statistics
-------------------------------------------------------------------------------
Message Type            Received      Transmitted      Forwarded
-------------------------------------------------------------------------------
```

```
General Queries         0             0               0
Group Queries           0             0               0
Group-Source Queries    0             0               0
V1 Reports              0             0               0
V2 Reports              0             0               0
V3 Reports              0             0               0
V2 Leaves               0             0               0
Unknown Type            0             N/A             0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length              : 0
Bad IP Checksum         : 0
Bad IGMP Checksum       : 0
Bad Encoding            : 0
No Router Alert         : 0
Zero Source IP          : 0

Send Query Cfg Drops    : 0
Import Policy Drops     : 0
Exceeded Max Num Groups : 0
===============================================================================
*A:ALA-48>show>service>id>snooping#
```

The following example applies to the 7450 ESS:

```
*A:ALA-42>show>service>id>igmp-snooping>snooping# all
===============================================================================
IGMP Snooping info for service 100
===============================================================================
IGMP Snooping Base info
-------------------------------------------------------------------------------
Admin State : Up
Querier     : 10.20.1.6 on SAP 2/2/1:100
-------------------------------------------------------------------------------
Sap/Sdp                  Oper    MRtr  Send      Max Num   MVR        Num
Id                       State   Port  Queries   Groups    From-VPLS  Groups
-------------------------------------------------------------------------------
sap:1/1/4                Up      No    Disabled  No Limit  Local      1
sap:2/2/1:100            Up      No    Disabled  No Limit  Local      0
-------------------------------------------------------------------------------
IGMP Snooping Querier info
-------------------------------------------------------------------------------
Sap Id                  : 2/2/1:100
IP Address              : 10.20.1.6
Expires                 : 3s
Up Time                 : 0d 00:15:23
Version                 : 3

General Query Interval  : 2s
Query Response Interval : 1.0s
Robust Count            : 2
-------------------------------------------------------------------------------
IGMP Snooping Multicast Routers
-------------------------------------------------------------------------------
MRouter          Sap/Sdp Id             Up Time        Expires   Version
-------------------------------------------------------------------------------
```

```
10.20.1.6        2/2/1:100              0d 00:15:24    2s        3
-------------------------------------------------------------------------------
Number of mrouters: 1
-------------------------------------------------------------------------------
IGMP Snooping Proxy-reporting DB
-------------------------------------------------------------------------------
Group Address    Mode     Up Time        Num Sources
-------------------------------------------------------------------------------
239.0.0.0        exclude  0d 00:00:04    0
-------------------------------------------------------------------------------
Number of groups: 1
-------------------------------------------------------------------------------
IGMP Snooping SAP 1/1/4 Port-DB
-------------------------------------------------------------------------------
Group Address    Mode     Type     From-VPLS  Up Time        Expires  Num Src
-------------------------------------------------------------------------------
239.0.0.0        exclude  dynamic  local      0d 00:00:05    4s       0
-------------------------------------------------------------------------------
Number of groups: 1
-------------------------------------------------------------------------------
IGMP Snooping SAP 2/2/1:100 Port-DB
-------------------------------------------------------------------------------
Group Address    Mode     Type     From-VPLS  Up Time        Expires  Num Src
-------------------------------------------------------------------------------
Number of groups: 0
-------------------------------------------------------------------------------
IGMP Snooping Static Source Groups
-------------------------------------------------------------------------------
IGMP Snooping Statistics
-------------------------------------------------------------------------------
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         463           0             463
Group Queries           0             0             0
Group-Source Queries    0             0             0
V1 Reports              0             0             0
V2 Reports              0             0             0
V3 Reports              4             4             0
V2 Leaves               0             0             0
Unknown Type            0             N/A           0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length              : 0
Bad IP Checksum         : 0
Bad IGMP Checksum       : 0
Bad Encoding            : 0
No Router Alert         : 0
Zero Source IP          : 0

Send Query Cfg Drops    : 0
Import Policy Drops     : 0
Exceeded Max Num Groups : 0

MVR From VPLS Cfg Drops : 0
MVR To SAP Cfg Drops    : 0
-------------------------------------------------------------------------------
IGMP Snooping Multicast VPLS Registration info
-------------------------------------------------------------------------------
```

```
IGMP Snooping Admin State : Up

MVR Admin State           : Down
MVR Policy                : None
-------------------------------------------------------------------------------
Local SAPs/SDPs
-------------------------------------------------------------------------------
Svc Id     Sap/Sdp                Oper      From       Num Local
           Id                     State     VPLS       Groups
-------------------------------------------------------------------------------
100        sap:1/1/4              Up        Local      1
100        sap:2/2/1:100          Up        Local      0
-------------------------------------------------------------------------------
MVR SAPs (from-vpls=100)
-------------------------------------------------------------------------------
Svc Id     Sap/Sdp                Oper      From       Num MVR
           Id                     State     VPLS       Groups
-------------------------------------------------------------------------------
No MVR SAPs found.
-------------------------------------------------------------------------------
*A:ALA-42>show>service>id>snooping#
```

The following table describes the show all service-id command output fields:

Table 35 **IGMP Snooping Show All Service-ID Field Descriptions**

Label	Description
Admin State	The administrative state of the IGMP instance
Querier	The address of the IGMP querier on the IP subnet to which the interface is attached
Sap/Sdp Id	The SAP and SDP IDs of the service ID
Oper State	The operational state of the SAP and SDP IDs of the service ID
Mrtr Port	The multicast router port
Send Queries	Send-queries command is enabled or disabled
Max Num Groups	The maximum number of multicast groups that can be joined on this SAP or SDP
MVR From VPLS	MVR from VPLS enabled
Num Groups	The actual number of multicast groups that can be joined on this SAP or SDP

mrouters

Syntax **mrouters** [**detail**]

Context show>service>id>igmp-snooping

Description This command displays all multicast routers.

Parameters **detail** — Displays detailed information.

Output The following output is an example of IGMP snooping mrouter information.

Sample Output

```
*A:ALA-48# show service id 700 igmp-snooping mrouters
===============================================================================
IGMP Snooping Multicast Routers for service 700
===============================================================================
MRouter          Sap/Sdp Id             Up Time        Expires   Version
-------------------------------------------------------------------------------
Number of mrouters: 0
===============================================================================
*A:ALA-48#
```

The following table describes the show igmp-snooping mrouters output fields:

Table 36 **IGMP Snooping Mrouters Field Descriptions**

Label	Description
MRouter	The multicast router port
Sap/Sdp Id	The SAP and SDP ID multicast router ports
Up Time	The length of time the mrouter has been up
Expires	The amount of time left before the query interval expires
Version	The configured version of IGMP running on this interface

mvr

Syntax **mvr**

Context show>service>id>igmp-snooping

Description This command displays Multicast VPLS Registration (MVR) information.

Output The following output is an example of IGMP snooping mvr information.

Sample Output

```
A:ALA-1>show>service>id>snooping# mvr
===============================================================================
IGMP Snooping Multicast VPLS Registration info for service 10
===============================================================================
```

```
IGMP Snooping Admin State : Up

MVR Admin State           : Up
MVR Policy                : mvr-policy
-------------------------------------------------------------------------------
Local SAPs/SDPs
-------------------------------------------------------------------------------
Svc Id     Sap/Sdp                Oper      From       Num Local
           Id                     State     VPLS       Groups
-------------------------------------------------------------------------------
100         sap:1/1/10:10         Up       Local      100
100         sap:1/1/10:20         Up       Local      100
-------------------------------------------------------------------------------
MVR SAPs (from-vpls=10)
-------------------------------------------------------------------------------
Svc Id     Sap/Sdp                Oper      From       Num MVR
           Id                     State     VPLS       Groups
-------------------------------------------------------------------------------
20         sap:1/1/4:100          Up        10         100
30         sap:1/1/31:10.10       Up        10         100
===============================================================================
A:ALA-1>show>service>id>snooping#
```

The following table describes the show igmp-snooping mvr output fields:

Table 37 **IGMP Snooping Mvr Field Descriptions**

Label	Description
IGMP Snooping Admin State	The IGMP snooping administrative state
MVR Admin State	The MVR administrative state
MVR Policy	The MVR policy name
Svc ID	The service ID
Sap/SDP	The SAP or SDP ID
Oper State	The operational state
From VPLS	The originating VPLS name
Num Local Groups	The number of local groups

port-db

Syntax **port-db sap** *sap-id* [**detail**]
port-db sap *sap-id* **group** *grp-address*
port-db sdp-*id*:*vc-id* [**detail**]
port-db sdp *sdp-id*:*vc-id* **group** *grp-address*

Context show>service>id>igmp-snooping

Description This command displays information on the IGMP snooping port database for the VPLS service.

Parameters *grp-ip-address* — Displays the IGMP snooping port database for a specific multicast group address.

sap-id — Displays the IGMP snooping port database for a specific SAP.

sdp-id — Displays only IGMP snooping entries associated with the specified mesh SDP or spoke SDP. For a spoke SDP, the VC ID must be specified; for a mesh SDP, the VC ID is optional.

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID for which to display information.

Values 1 to 4294967295

Default For mesh SDPs only, all VC IDs

grp-address — Displays IGMP snooping statistics matching the specified group address.

ip-address — Displays IGMP snooping statistics matching one particular source within the multicast group.

Output The following output is an example of IGMP snooping port database information.

Sample Output

The following example applies to the 7450 ESS:

```
A:ALA-1>show>service>id>snooping# port-db sap 1/1/2
===============================================================================
IGMP Snooping SAP 1/1/2 Port-DB for service 10
===============================================================================
Group Address    Mode     Type     From-VPLS  Up Time        Expires  Num Src
-------------------------------------------------------------------------------
239.0.0.1        include  dynamic  Local      0d 00:04:44    0s       2
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================
```

The following example applies to the 7750 SR:

```
A:ALA-1>show>service>id>snooping# port-db sap 1/1/2
===============================================================================
IGMP Snooping SAP 1/1/2 Port-DB for service 10
===============================================================================
Group Address    Mode     Type     Up Time        Expires   Num Sources
-------------------------------------------------------------------------------
239.0.0.1        include  dynamic  0d 00:04:44    0s        2
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================
```

```
A:ALA-1>show>service>id>snooping# port-db sap 1/1/2 detail
===============================================================================
IGMP Snooping SAP 1/1/2 Port-DB for service 10
===============================================================================
IGMP Group 239.0.0.1
-------------------------------------------------------------------------------
Mode             : include              Type             : dynamic
Up Time          : 0d 00:04:57          Expires          : 0s
Compat Mode      : IGMP Version 3
V1 Host Expires  : 0s                   V2 Host Expires  : 0s
-------------------------------------------------------
Source Address  Up Time      Expires  Type     Fwd/Blk
-------------------------------------------------------
10.1.1.1         0d 00:04:57  20s      dynamic  Fwd
10.1.1.2         0d 00:04:57  20s      dynamic  Fwd
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================
A:ALA-1>show>service>id>snooping#
```

proxy-db

Syntax **proxy-db** [**detail**]
proxy-db group *grp-address*

Context show>service>id>igmp-snooping

Description This command displays information on the IGMP snooping proxy reporting database for the VPLS service.

Parameters *grp-ip-address* — Displays the IGMP snooping proxy reporting database for a specific multicast group address.

Output The following output displays an example of IGMP snooping proxy database information.

Sample Output

```
A:ALA-1>show>service>id>snooping# proxy-db
===============================================================================
IGMP Snooping Proxy-reporting DB for service 10
===============================================================================
Group Address    Mode     Up Time        Num Sources
-------------------------------------------------------------------------------
239.0.0.1        include  0d 00:05:40    2
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================

A:ALA-1>show>service>id>snooping# proxy-db detail
===============================================================================
IGMP Snooping Proxy-reporting DB for service 10
===============================================================================
-------------------------------------------------------------------------------
```

```
IGMP Group 239.0.0.1
-------------------------------------------------------------------------------
Up Time : 0d 00:05:54                  Mode : include
------------------------------
Source Address  Up Time
------------------------------
10.1.1.1         0d 00:05:54
10.1.1.2         0d 00:05:54
-------------------------------------------------------------------------------
Number of groups: 1
===============================================================================
A:ALA-1>show>service>id>snooping#
```

querier

Syntax **querier**

Context show>service>id>igmp-snooping

Description This command displays information on the IGMP snooping queriers for the VPLS service.

Output The following output displays an example of IGMP snooping querier information.

Sample Output

```
A:ALA-1>show>service>id>snooping# querier
===============================================================================
IGMP Snooping Querier info for service 10
===============================================================================
Sap Id                  : 1/1/1
IP Address              : 10.10.10.1
Expires                 : 6s
Up Time                 : 0d 00:56:50
Version                 : 3

General Query Interval  : 5s
Query Response Interval : 2.0s
Robust Count            : 2
===============================================================================
A:ALA-1>show>service>id>snooping#
```

static

Syntax **static** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context show>service>id>igmp-snooping

Description This command displays information on static IGMP snooping source groups for the VPLS service.

Parameters *sap-id* — Displays static IGMP snooping source groups for a specific SAP.

sdp-id — Displays the IGMP snooping source groups for a specific spoke or mesh SDP.

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID for which to display information.

Values 1 to 4294967295

Default For mesh SDPs only, all VC IDs

Output The following output displays an example of static IGMP snooping information.

Sample Output

```
A:ALA-1>show>service>id>snooping# static
===============================================================================
IGMP Snooping Static Source Groups for SAP 1/1/2
-------------------------------------------------------------------------------
Source           Group
-------------------------------------------------------------------------------
*                239.0.0.2
*                239.0.0.3
-------------------------------------------------------------------------------
Static (*,G)/(S,G) entries: 2

-------------------------------------------------------------------------------
IGMP Snooping Static Source Groups for SDP 10:10
-------------------------------------------------------------------------------
Source           Group
-------------------------------------------------------------------------------
10.1.1.1          239.0.0.10
-------------------------------------------------------------------------------
Static (*,G)/(S,G) entries: 1
===============================================================================
A:ALA-1>show>service>id>snooping#
```

statistics

Syntax **statistics** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context show>service>id>igmp-snooping

Description This command displays IGMP snooping statistics for the VPLS service.

Parameters *sap-id* — Displays IGMP snooping statistics for a specific SAP.

sdp-id — Displays the IGMP snooping statistics for a specific spoke or mesh SDP.

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID for which to display information.

Values 1 to 4294967295

Default For mesh SDPs only, all VC IDs

Output The following output displays an example of IGMP snooping statistics information.

Sample Output

```
A:ALA-1>show>service>id>snooping# statistics
===============================================================================
IGMP Snooping Statistics for service 1
===============================================================================
Message Type            Received      Transmitted   Forwarded
-------------------------------------------------------------------------------
General Queries         4             0             4
Group Queries           0             0             0
Group-Source Queries    0             0             0
V1 Reports              0             0             0
V2 Reports              0             0             0
V3 Reports              0             0             0
V2 Leaves               0             0             0
Unknown Type            0             N/A           0
-------------------------------------------------------------------------------
Drop Statistics
-------------------------------------------------------------------------------
Bad Length              : 0
Bad IP Checksum         : 0
Bad IGMP Checksum       : 0
Bad Encoding            : 0
No Router Alert         : 0
Zero Source IP          : 0

Send Query Cfg Drops    : 0
Import Policy Drops     : 0
Exceeded Max Num Groups : 0

MVR From VPLS Cfg Drops : 0
MVR To SAP Cfg Drops    : 0
===============================================================================
A:ALA-1>show>service>id>snooping#
```

mc-ecmp-balance

Syntax **mc-ecmp-balance** [**detail**]

Context show>router>pim

Description This command displays multicast balance information.

Parameters **detail** — Displays detailed information.

Output The following output displays an example of PIM mc-ecmp-balance information.

Sample Output

```
A:ALA-48>config>router>pim# show router pim mc-ecmp-balance
===============================================================================
PIM ECMP Balance
===============================================================================
MC-ECMP-Balance                  : Disabled
Rebalance in progress            : No
```

```
Last Rebalance Time               : 11/13/2007 09:03:10
Rebalance Type                    : Unknown
Optional Threshold Used           : 0
Mc Ecmp Balance Hold Time         : None
===============================================================================
A:ALA-48>config>router>pim#
```

group

Syntax **group** [*grp-ip-address*]
group summary

Context show>router>igmp

Description This command displays IGMP group information.

group-interface

Syntax **group-interface** [**fwd-service** *service-id*] [*ip-int-name*] [**detail**]

Context show>router>igmp

Description This command displays IGMP group-interface information.

hosts

Syntax **hosts** [**group** *grp-address*] [**detail**] [**fwd-service** *service-id*] [**grp-interface** *ip-int-name*]
hosts [**host** *ip-address*] [**group** *grp-address*] [**detail**]
hosts summary

Context show>router>igmp

Description This command displays IGMP hosts information.

interface

Syntax **interface** [*ip-int-name* | *ip-address*] [**group**] [*grp-ip-address*] [**detail**]

Context show>router>igmp

Description This command displays IGMP interface information.

mcast-reporting-statistics

Syntax **mcast-reporting-statistics** [*ip-address*]

Context show>router>igmp

Description This command displays IGMP mcast reporting statistics.

ssm-translate

Syntax **ssm-translate** [*interface-name*]

Context show>router>igmp

Description This command displays SSM translate configuration information.

static

Syntax **static** [*ip-int-name* | *ip-addr*]

Context show>router>igmp

Description This command displays IGMP static group/source configuration information.

statistics

Syntax **statistics** [*ip-int-name* | *ip-address*]
statistics host [*ip-address*]

Context show>router>igmp

Description This command displays IGMP statistics information.

status

Syntax **status**

Context show>router>igmp

Description This command displays IGMP status information.

tunnel-interface

Syntax **tunnel-interface**

Context show>router>igmp

Description This command displays IGMP tunnel-interface information.

2.11.2.4 Clear Commands

id

Syntax **id** *service-id*

Context clear>service
clear>service>statistics

Description This command clears the identification for a specific service.

Parameters *service-id* — The ID that uniquely identifies a service.

Values service-id: 1 to 214748364
svc-name: A string up to 64 characters.

arp-host

Syntax **arp-host** [**wholesaler** *service-id*] [**sap** *sap-id* | **interface** *interface-name* | **ip-address** *ip-address*[*/mask*] | **mac** *ieee-address* | {[**port** *port-id*] [**no-inter-dest-id** | **inter-dest-id** *inter-dest-id*]}] [**detail**]
arp-host statistics [**sap sap-id** | **interface** *interface-name*]
arp-host summary [**interface** *interface-name* | **saps**]

Context clear>service>id

Description This command clears ARP host data.

authentication

Syntax **authentication**

Context clear>service>id

Description This command enables the context to clear session authentication information.

statistics

Syntax **statistics**

Context clear>service>id>authentication

Description This command clears session authentication statistics for this service.

igmp-snooping

Syntax **igmp-snooping**

Context show>service>id

Description This command displays the Internet Group Management Protocol (IGMP) snooping information.

port-db

Syntax **port-db** {**sap** *sap-id* | **sdp** *sdp-id:vc-id*} [**group** *grp-address* [**source** *ip-address*]]

Context clear>service>id>igmp-snooping

Description This command clears the information on the IGMP snooping port database for the VPLS service.

Parameters *sap-id* — Clears IGMP snooping statistics matching the specified SAP ID and optional encapsulation value.

sdp-id — Clears only IGMP snooping entries associated with the specified mesh SDP or spoke SDP. For a spoke SDP, the VC ID must be specified; for a mesh SDP, the VC ID is optional.

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID for which to clear information.

Values 1 to 4294967295

Default For mesh SDPs only, all VC IDs

grp-address — Clears IGMP snooping statistics matching the specified group address.

ip-address — Clears IGMP snooping statistics matching one particular source within the multicast group.

querier

Syntax **querier**

Context clear>service>id>igmp-snooping

Description This command clears the information on the IGMP snooping queriers for the VPLS service.

statistics

Syntax **statistics** [**sap** *sap-id* | **sdp** *sdp-id:vc-id*]

Context clear>service>id>mld-snooping

Description This command clears MLD snooping statistics.

Parameters *sap-id* — Clears all MLD snooping SAP statistics.

sdp-id:vc-id — Clears all MLD snooping SDP statistics.

msap

Syntax **msap** {*msap-id* | **idle-only**}

Context clear>service>id

Description This command clears Managed SAP (MSAP) information.

This command can remove an MSAP with active subscribers still associated with the MSAP. You can specify the **idle-only** parameter to clear only MSAPs in an idle state; MSAPs with active subscribers are not cleared.

Parameters *msap-id* — Specifies a Managed SAP ID.

Values

dot1q	[port-id | lag-id]:qtag1
qinq	[port-id | lag-id]:qtag1.qtag2
qtag1	0 to 4094
qtag2	0 to 4094

idle-only — Specifies to clear only MSAPs in an idle state.

msap-policy

Syntax **msap-policy** *msap-policy-name*

Context clear>service>id

Description This command clears Managed SAPs created by the Managed SAP policy.

Parameters *msap-policy-name* — Clears an existing MSAP policy.

detected-protocols

Syntax **detected-protocols** {**all** | **sap** *sap-id* | **spoke-sdp** *sdp-id*[*:vc-id*]}

Context clear>service>id>stp

Description RSTP automatically falls back to STP mode when it receives an STP BPDU. The **clear detected-protocols** command forces the system to revert to the default RSTP mode on the SAP or spoke SDP.

Parameters **all** — Clears all detected protocol information.

sap-id — Clears the specified lease state SAP information.

sdp-id — The SDP ID to be cleared.

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID to be cleared.

Values 1 to 4294967295

statistics

Syntax **statistics** {**all** | **sap** *sap-id* | **sdp** *sdp-id:vc-id*}]

Context clear>service>id>igmp-snooping

Description This command clears IGMP snooping statistics for the VPLS service.

Parameters *sap-id* — Clears IGMP snooping statistics matching the specified SAP ID and optional encapsulation value.

sdp-id — Clears only IGMP snooping entries associated with the specified mesh SDP or spoke SDP. For a spoke SDP, the VC ID must be specified, for a mesh SDP, the VC ID is optional.

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID for which to clear statistics.

Values 1 to 4294967295

Default For mesh SDPs only, all VC IDs

mfib

Syntax **mfib**

Context clear>service>id

Description This command enables the context to clear multicast FIB info for the VPLS service.

statistics

Syntax **statistics** {**all** | **group** *grp-address*}

Context clear>service>id>mfib

Description This command clears multicast FIB statistics for the VPLS service.

Parameters *grp-address* — Specifies an IGMP multicast group address that receives data on an interface.

mld-snooping

Syntax **mld-snooping**

Context clear>service>id

Description This command enables the context to clear MLD snooping-related data.

port-db

Syntax **port-db sap** *sap-id* [**group** *grp-ipv6-address*]
port-db sap *sap-id* **group** *grp-ipv6-address* **source** *src-ipv6-address*
port-db sdp *sdp-id*:*vc-id* [**group** *grp-ipv6-address*]
port-db sdp *sdp-id*:*vc-id* **group** *grp-ipv6-address* **source** *src-ipv6-address*

Context clear>service>id>mld-snooping

Description This command clears MLD snooping port-db group data.

querier

Syntax **querier**

Context clear>service>id>mld-snooping

Description This command clears MLD snooping querier information.

statistics

Syntax **querier**

Context clear>service>id>mld-snooping

Description This command clears MLD statistics.

arp

Syntax **arp** {**all** | *ip-address*}
arp interface [*ip-int-name* | *ip-address*]

Context clear>router

Description This command clears all or specific ARP entries. The scope of ARP cache entries cleared depends on the command line option(s) specified.

Parameters **all** — Clears all ARP cache entries.

ip-address — Clears the ARP cache entry for the specified IP address.

ip-int-name — Clears all ARP cache entries for the IP interface with the specified name.

ip-addr — Clears all ARP cache entries for the specified IP interface with the specified IP address.

2.11.2.5 Debug Commands

id

Syntax [**no**] **id** *service-id*

Context debug>service>id

Description This command enables debugging for the specified service ID.

The **no** form of this command disables the debugging

arp-host

Syntax [**no**] **arp-host**

Context debug>service>id

Description This command enables ARP host debugging.

The **no** form of this command disables ARP host debugging

mld-snooping

Syntax [**no**] **mld-snooping**

Context debug>service>id

Description This command enables and configures MLD-snooping debugging.

The **no** form of this command disables MLD-snooping debugging

detail-level

Syntax **detail-level** {**low** | **medium** | **high**}
no detail-level

Context debug>service>id>mld

Description This command enables and configures the MLD tracing detail level.

The **no** form of this command disables the MLD tracing detail level.

mac

Syntax [**no**] **mac** *ieee-address*

Context debug>service>id>mld

Description This command shows MLD packets for the specified MAC address.

The **no** form of this command disables the MAC debugging.

mode

Syntax **mode** {**dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}
no mode

Context debug>service>id>mld

Description This command enables and configures the MLD tracing mode.

The **no** form of this command disables the configures the MLD tracing mode.

sap

Syntax [**no**] **sap** *sap-id*

Context debug>service>id>mld

Description This command shows MLD packets for a specific SAP.

The **no** form of this command disables the debugging for the SAP.

sdp

Syntax [**no**] **sdp** *sdp-id:vc-id*

Context debug>service>id>mld

Description This command shows MLD packets for a specific SDP.

The **no** form of this command disables the debugging for the SDP.

2.11.2.6 Tools Commands

msap

Syntax **msap** [**service** *service-id*]

Context tools>dump>service

Description This command provides MSAP-related statistics for the chassis and tracks the number of sticky MSAPs. The statistics include the total number of MSAPs, sticky MSAPs, and idle MSAPs. When a subscriber disconnects from a sticky MSAP, it transitions to an idle MSAP. The idle MSAP transitions back to a sticky MSAP when the subscriber reconnects. A large number of idle MSAPs during peak network hours indicate that an idle MSAP cleanup is required. The total MSAP statistic counts both the total number of traditional MSAPs (non-sticky) and total number of sticky MSAPs. Idle MSAPs are counted towards the total number of sticky MSAPs.

Parameters *service-id* — Specifies the service for which to show the MSAP-related statistics. The service is indicated in the chassis.

3 DHCP Management

3.1 DHCP Principles

In a Triple Play network, client devices (such as a routed home gateway, a session initiation protocol (SIP) phone, or a set-top box) use the Dynamic Host Configuration Protocol (DHCP) to dynamically obtain their IP address and other network configuration information.

DHCP is defined and shaped by several RFCs and drafts in the IETF DHCP working group including the following:

- RFC 1534, *Interoperation Between DHCP and BOOTP*
- RFC 2131, *Dynamic Host Configuration Protocol*
- RFC 2132, *DHCP Options and BOOTP Vendor Extensions*
- RFC 3046, *DHCP Relay Agent Information Option*

The DHCP operation is shown in Figure 14.

Figure 14 **IP Address Assignment with DHCP**

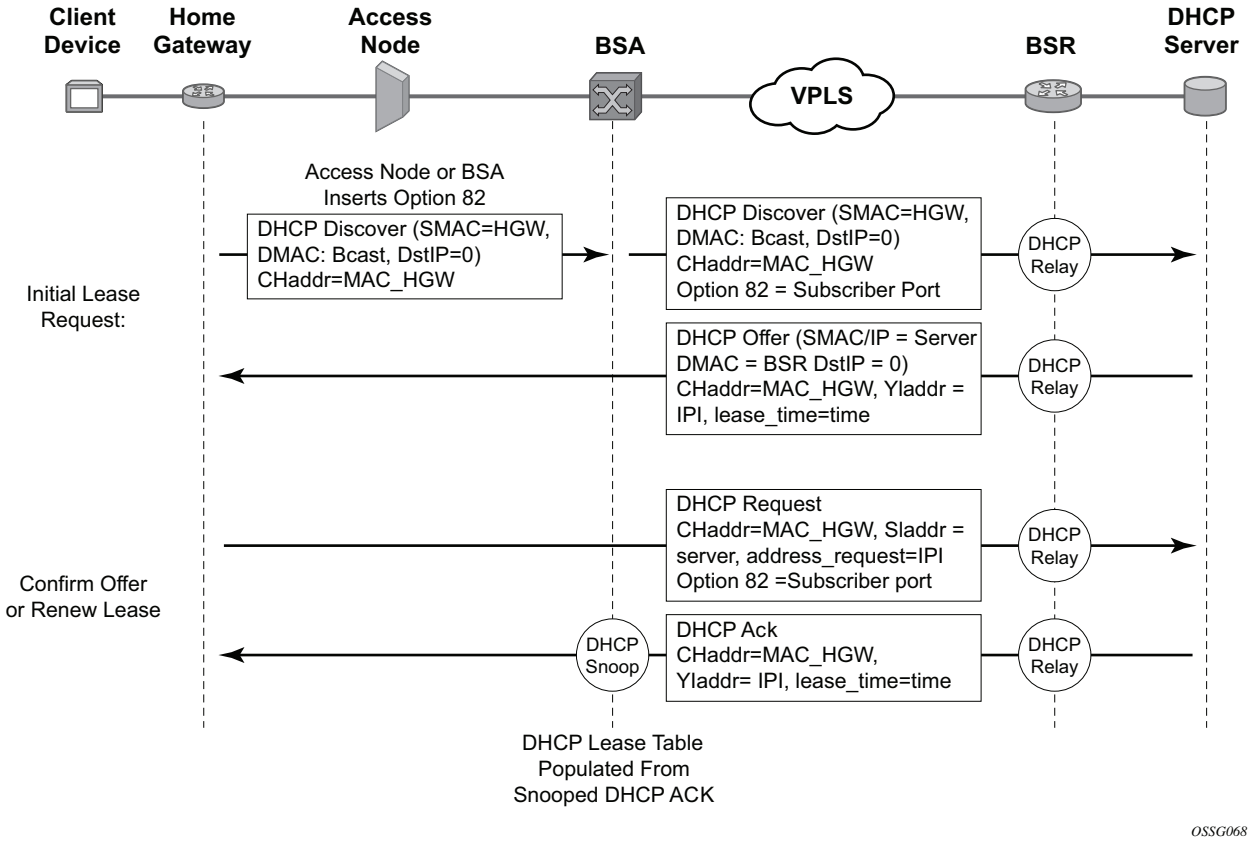

1. During bootup, the client device sends a DHCP discover message to get an IP address from the DHCP Server. The message contains:
   - Destination MAC address — broadcast
   - Source MAC address — the MAC address of the client device
   - Client hardware address — the MAC address of the client device

   If this message passes through a DSLAM or other access node, typically the relay information option (Option 82) field is added, indicating shelf, slot, port, VPI, VCI, and so on, to identify the subscriber.

   DHCP relay is enabled on the first IP interface in the upstream direction. Depending on the scenario, the DSLAM, BSA, or the BSR relays the discover message as a unicast packet toward the configured DHCP server. DHCP relay is configured to insert the GIADDR (gateway IP address) in order to indicate to the DHCP server in which subnet an address should be allocated.

2. The DHCP server looks up the client MAC address and Option 82 information in its database. If the client is recognized and authorized to access the network, an IP address is assigned and a DHCP offer message returned. The BSA or BSR then relays this back to the client device.
3. It is possible that the discover reached more than one DHCP server and more than one offer is returned. The client selects one of the offered IP addresses and confirms that it wants to use this in a DHCP request message, sent as unicast to the DHCP server that offered it.
4. The DHCP server confirms that the IP address is still available, updates its database to indicate it is now in use, and replies with a DHCP ACK message back to the client. The ACK also contains the Lease Time of the IP address.

3.2 DHCP Features

3.2.1 DHCP Relay

Since DHCP requests are broadcast packets that normally do not propagate outside of their IP subnet, a DHCP relay agent intercepts such requests and forwards them as unicast messages to a configured DHCP server.

When forwarding a DHCP message, the relay agent sets the GIADDR in the packet to the IP address of its ingress interface. This allows DHCP clients to use a DHCP server on a remote network. From both a scalability and a security point of view, it is recommended that the DHCP relay agent is positioned as close as possible to the client terminals.

DHCP relay is used in a Layer 3 environment, and thus is only supported in IES services and VPRN services. VPRN is only supported on the 7750 SR.

When DHCP clients and servers are in different VPRN routing instances of which one is the Base routing instance, route leaking (GRT-leaking) should be used to relay DHCPv4 and DHCPv6 messages between a VPRN and the Global Routing Table (GRT).

While DHCP relay is not implemented in a VPLS, it is still possible to insert or modify Option 82 information.

In a Routed CO environment in the 7750 SR, the subscriber interface's group interface DHCP relay is stateful.

3.2.2 DHCPv4 Relay Proxy

In network deployments where DHCPv4 client subnets cannot be leaked in the DHCPv4 server routing instance, unicast renewal messages (DHCP ACKs) cannot be routed in the DHCPv4 server routing instance, as shown in Figure 15. The DHCP server sets the destination IP address of the DHCP ACK to the client IP address (ciaddr) as received in the DHCP REQUEST message. Because there is no route available for the client subnet in the DHCP server routing instance, the DHCP ACK cannot be delivered.

Figure 15 **Unicast Renewal Routing Problem**

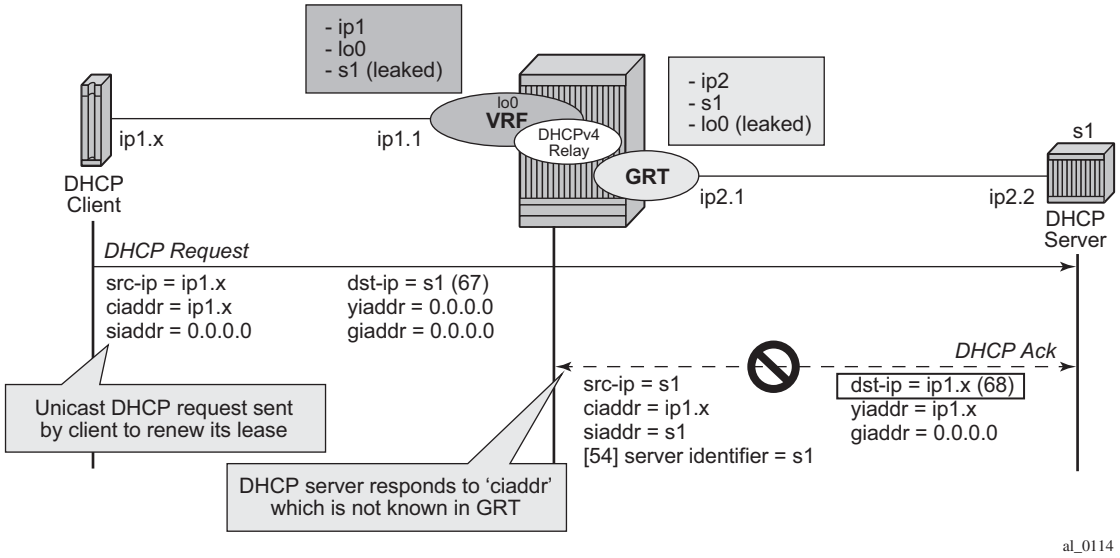

The unicast renewal routing problem shown in Figure 15 can be solved with a relay proxy function that enhances the DHCPv4 relay. With the **relay-proxy** command in the DHCPv4 relay on a regular interface or group interface, the unicast renewals are now also relayed to the DHCPv4 server, as described below and shown in Figure 16:

- In the client to server direction, the source IP address is updated and the gateway IP address (gi-address) field is added before sending the message to the intended DHCP server (the message is not broadcasted to all configured DHCP servers.
- In the server to client direction, the GI address field is removed and the destination IP address is updated with the value of the IP address (yiaddr) field.

When **relay-proxy** is enabled, the GI address can be configured to any local address that is configured in the same routing instance. The GI address is the only address that must be leaked in the DHCPv4 server routing instance because a DHCPv4 server always sends the response on a relayed packet to the relay agent using the gi-address as the destination IP address.

By default, unicast DHCPv4 RELEASE messages are forwarded transparently by a relay proxy function. The optional **release-update-src-ip** flag updates the source IP address with the value that is used for all relayed DHCPv4 messages, as shown in Figure 16.

DHCPv4 FORCERENEW messages that are sent from a trusted external DHCPv4 server to a DHCPv4 relay agent configured as a relay proxy are forwarded to the DHCP client, if a corresponding DHCPv4 lease exists; otherwise, the DHCPv4 FORCERENEW messages are dropped.

Figure 16 **Relay Unicast Messages**

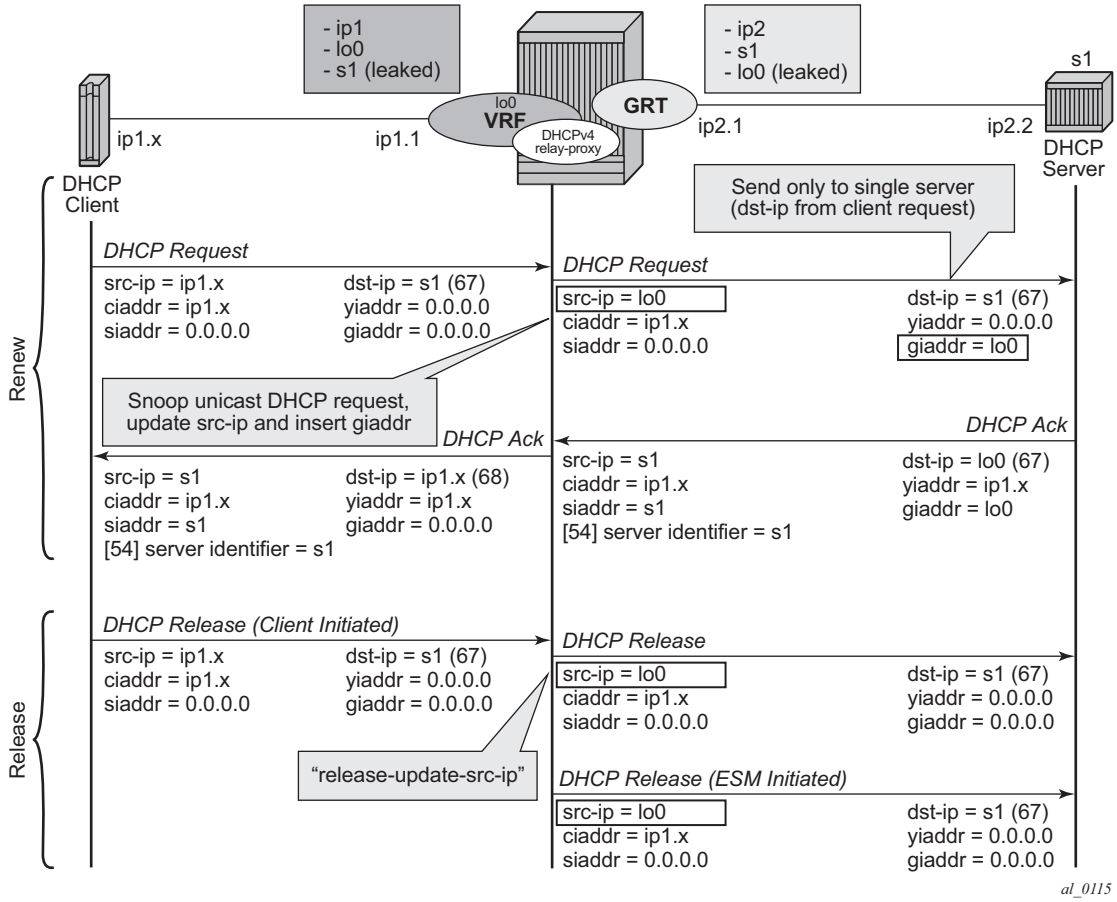

The **relay-proxy** command can also be used to hide the DHCPv4 server address for DHCP clients. This prevents the client from learning the DHCPv4 server infrastructure details such as the IP address and number of servers. Hiding infrastructure details helps in Denial of Service (DoS) prevention.

The optional **siaddr-override** *ip-address* parameter in **relay-proxy** enables DHCPv4 server IP address hiding toward the client. The client interacts with the relay proxy as if it is the DHCP server. In addition to the relay proxy functions as described earlier, the following actions are performed when DHCPv4 server IP address hiding is configured:

- In all DHCP messages to the client, the value of the following header fields and DHCP options containing the DHCP server IP address is replaced with the configured *ip-address*:
  - the source IP address

  - the siaddr field in the DHCPv4 header if it is not equal to zero in the message received from the server
  - the Server Identification option (DHCPv4 option 54) if present in the original server message
- The DHCP OFFER selection during initial binding. Only the first DHCP OFFER message is forwarded to the client. Subsequent DHCP OFFER messages from different servers are silently dropped.

The **siaddr-override** *ip-address* parameter can be any local address in the same routing instance. If DHCP relay lease split is enabled, **siaddr-override** *ip-address* has priority over the **emulated-server** *ip-address* configured in the proxy-server and is used as the source IP address.

The active DHCPv4 server IP address obtained from the DHCP OFFER selection is required for the IP address hiding function and is stored in the lease state record. Therefore **lease-populate** must be enabled on the interface when **siaddr-override** *ip-address* is configured.

Figure 17 shows the initial lease binding phase of a relay proxy with DHCP server address hiding enabled. In the absence of a DHCP lease state in the initial lease binding phase, the DHCP server IP address resulting from the OFFER selection is stored in a DHCP transaction cache. After successful lease binding, the DHCP server IP address is added to the lease state record.

In a host creation failure scenario, if no transaction cache or lease state is available when a DHCP REQUEST message is received, then the DHCP REQUEST is silently dropped. The drop reason can be found by enabling DHCP debug.

Figure 17 **DHCP Server IP Address Hiding/Initial Binding**

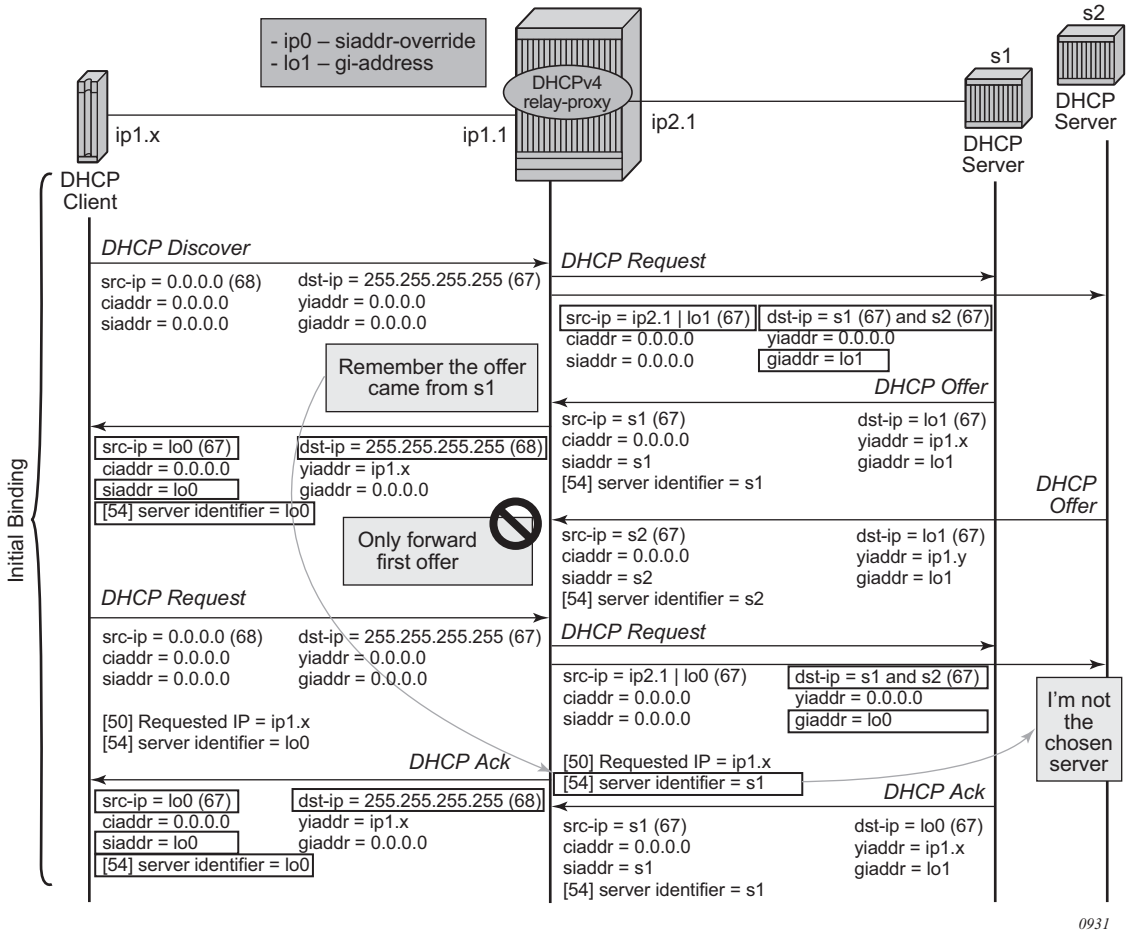

Figure 18 shows the lease renewal phase of a relay proxy with DHCP server address hiding enabled. A unicast REQUEST (renew) is relayed only to the DHCP server owning the lease. A broadcast REQUEST (rebind) is relayed to all configured DHCP servers.

During lease renewal, the DHCP server IP address can be updated in the lease state if the DHCP ACK is received from a different server. This optimizes the DHCP proxy relay operation in a DHCP server failover scenario. This is shown in Figure 19.

Figure 18 **DHCP Server IP Address Hiding/Lease Renewal**

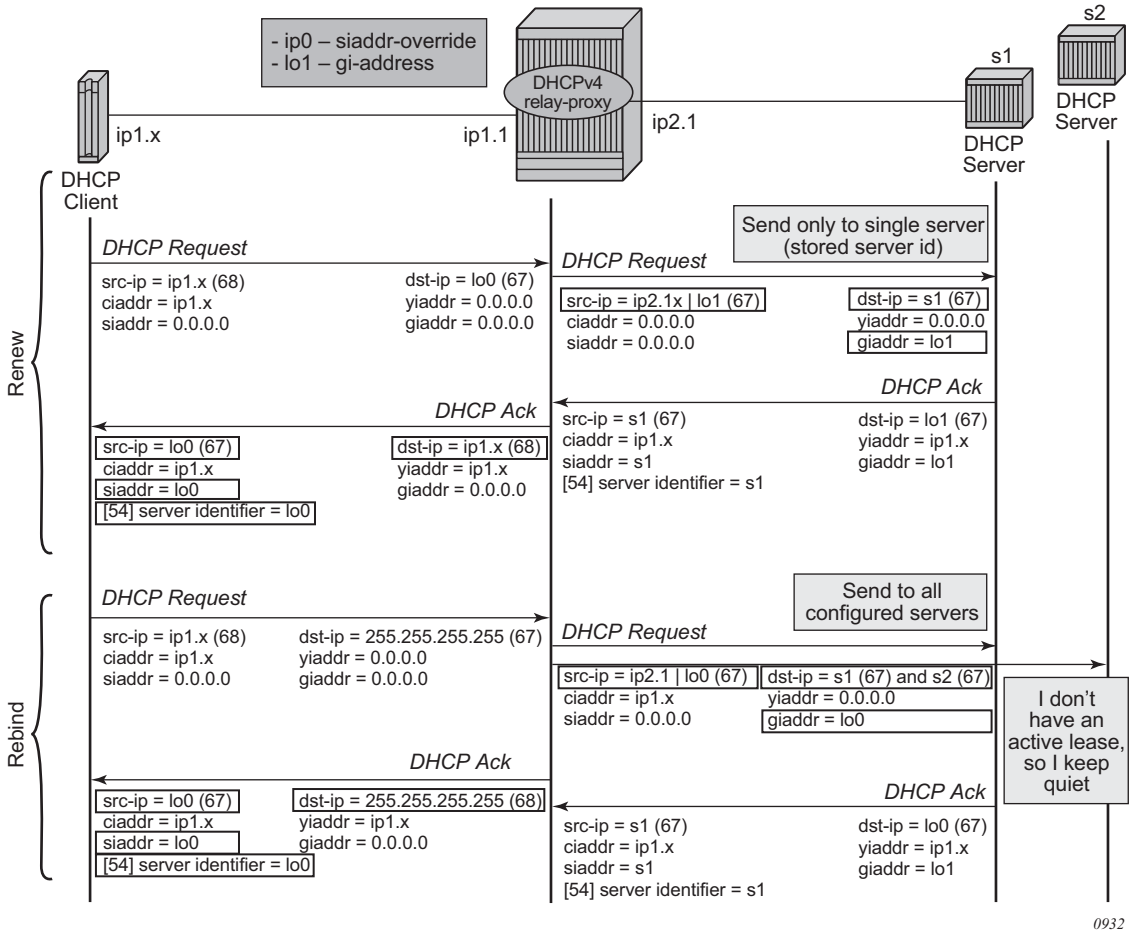

Figure 19 **DHCP Server IP Address Hiding, Lease Renewal With Active Server Failure**

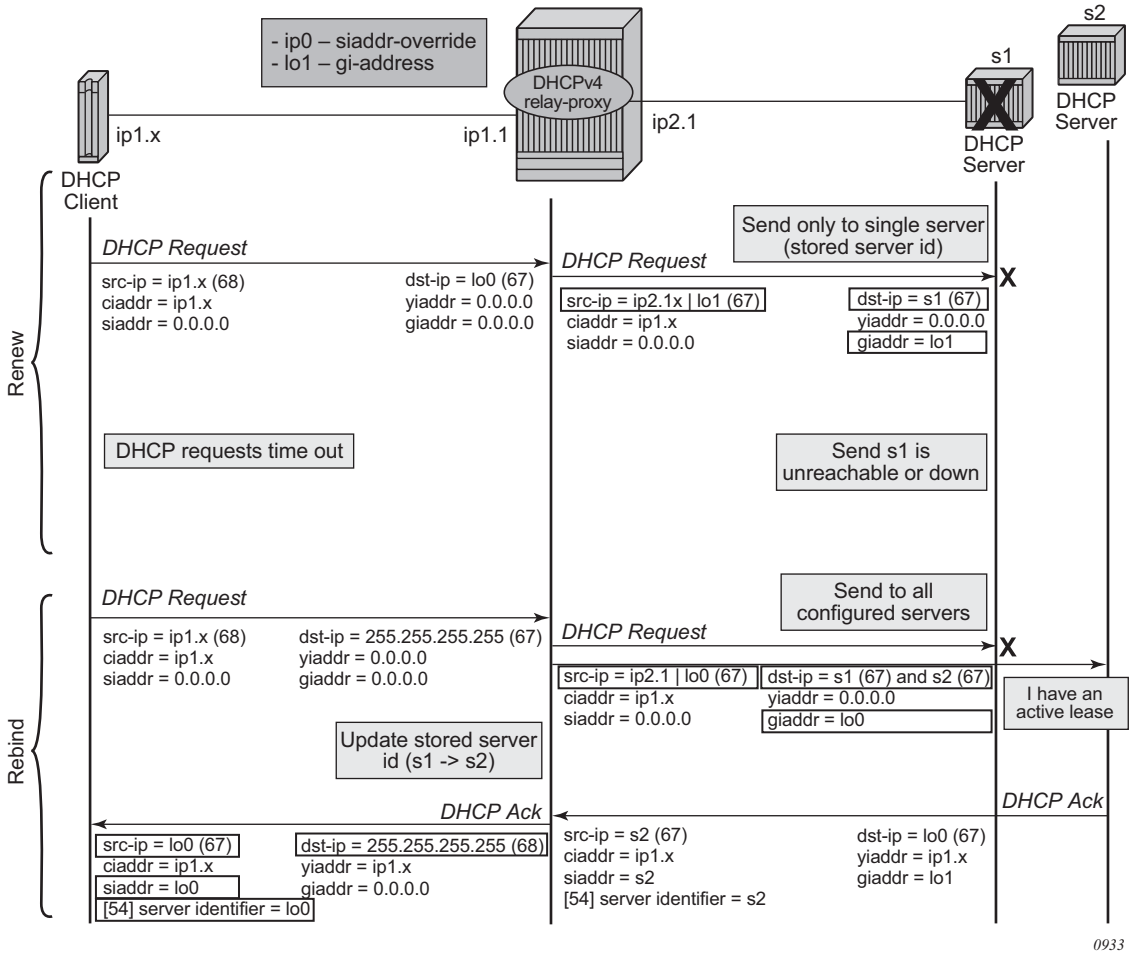

Figure 20 shows the release in a relay proxy scenario with DHCP server address hiding enabled. The RELEASE message is sent only to the DHCP server owning the lease. Optionally, the source IP address can be updated.

Figure 20 **DHCP Server IP Address Hiding, Release**

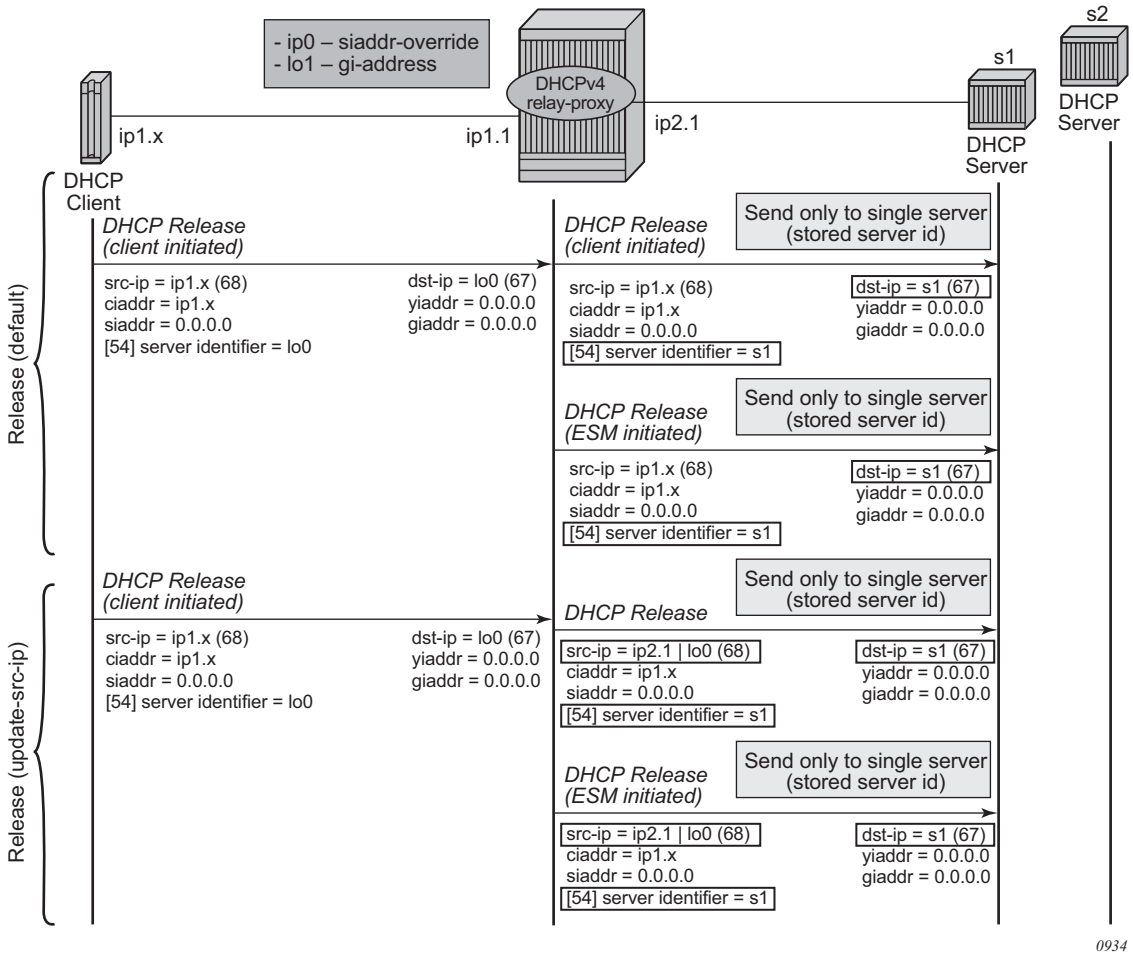

Relay proxy can be enabled on subscriber group-interfaces and regular interfaces in an IES or VPRN service.

For retail subscriber interfaces, **relay-proxy** is configured at the subscriber-interface dhcp CLI context, as shown in the example that follows.

A relay proxy function is not supported with a double DHCPv4 relay (Layer 3 DHCPv4 relay in front of a 7750 DHCPv4 relay with **relay-proxy** enabled).

Configuration example

```
config>service>vprn
        interface "lo0" create
            address 192.0.2.10/32
            loopback
```

```
exit
interface "lo1" create
    address 192.0.2.11/32
    loopback
exit
subscriber-interface "sub-int-1" create
    address 10.1.0.254/24
    group-interface "group-int-1-1" create
        dhcp
            server 172.16.1.1
            lease-populate 32767
            relay-proxy release-update-src-ip siaddr-override 192.0.2.10
            gi-address 192.0.2.11 src-ip-addr
            no shutdown
         exit
    exit
exit
```

3.2.3 Subscriber Identification Using Option 82 Field

Option 82, or the relay information option is specified in RFC 3046, *DHCP Relay Agent Information Option*, and allows the router to append some information to the DHCP request that identifies where the original DHCP request came from.

There are two sub-options under Option 82:

- *Agent Circuit ID Sub-option* (RFC 3046, section 3.1): This sub-option specifies data which must be unique to the box that is relaying the circuit.
- *Remote ID Sub-option* (RFC 3046 section 3.2): This sub-option identifies the host at the other end of the circuit. This value must be globally unique.

Both sub-options are supported and can be used separately or together.

Inserting Option 82 information is supported independently of DHCP relay. However in a VPLS (when DHCP Relay is not configured), DHCP snooping must be enabled on the SAP to be able to insert Option 82 information.

When the **circuit id sub-option** field is inserted, it can take following values:

- *sap-id:* The SAP index (only under an IES or VPRN service)
- *ifindex:* The index of the IP interface (only under an IES or VPRN service)
- *ascii-tuple:* An ASCII-encoded concatenated tuple, consisting of [*system-name*|*service-id*|*interface-name*] (for VPRN or IES) or [*system-name*|*service-id*|*sap-id*] (for VPLS).
- *vlan-ascii-tuple*: An ASCII-encoded concatenated tuple, consisting of the ascii-tuple followed by Dot1p bits and Dot1q tags.

For VPRN, the ifindex is unique only within a VRF. The DHCP relay function automatically prepends the VRF ID to the ifindex before relaying a DHCP Request. VPRN is supported on the 7750 SR only.

When a DHCP packet is received with Option 82 information already present, the system can do one of three things. The available actions are:

- Replace — On ingress the existing information-option is replaced with the **information-option** parameter configured. On egress (towards the customer) the information option is stripped (per the RFC).
- Drop — The DHCP packet is dropped and a counter is incremented.
- Keep — The existing information is kept on the packet and the router does not add any additional information. On egress the information option is not stripped and is sent on to the downstream node.

In accordance with the RFC, the default behavior is to keep the existing information; except if the GIADDR of the packet received is identical to a local IP address on the router, then the packet is dropped and an error incremented regardless of the configured action.

The maximum packet size for a DHCP relay packet is 1500 bytes. If adding the Option 82 information would cause the packet to exceed this size, the DHCP client packet is discarded. This packet size limitation exists to ensure that there is no fragmentation on the end Ethernet segment where the DHCP server attaches.

In the downstream direction, the inserted Option 82 information should not be passed back towards the client (as per RFC 3046, *DHCP Relay Agent Information Option*). To enable downstream stripping of the Option 82 field, DHCP snooping should be enabled on the SDP or SAP connected to the DHCP server.

3.2.3.1 Trusted and Untrusted

There is a case where the relay agent can receive a request where the downstream node added Option 82 information without also adding a GIADDR (GIADDR of 0). In this case the default behavior is for the router to drop the DHCP request. This behavior is in line with the RFC.

The **trusted** command allows the router to forward the DHCP request even if it receives one with a GIADDR of 0 and Option 82 information attached. This could occur with older access equipment. In this case, the relay agent can modify the request's GIADDR to be equal to the ingress interface. This can occur when the action in the information option is **keep**, and the service is IES or VPRN. In the case where the Option 82 information gets replaced by the relay agent, either through explicit configuration or the VPLS DHCP Relay case, the original Option 82 information is lost, and the reason for enabling the trusted option is lost.

3.2.4 DHCP Snooping

This section discusses the Nokia routers acting as a Broadband Subscriber Aggregator (BSA) with Layer 2 aggregation towards a Broadband Subscriber Router (BSR).

A typical initial DHCP scenario is shown in Figure 21.

Figure 21 **Initial DHCP Scenario**

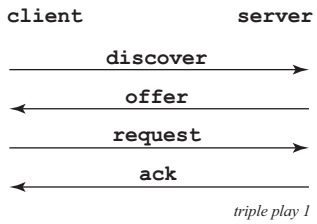

But, when the client already knows its IP address, it can skip the discover, as shown in Figure 22.

Figure 22 **DHCP Scenario with Known IP Address**

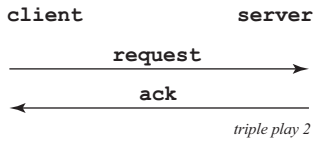

The BSA can copy packets designated to the standard UDP port for DHCP (port 67) to its control plane for inspection, this process is called DHCP snooping.

DHCP snooping can be performed in two directions:

1. From the client to the DHCP server (Discover or Request messages):

– to insert Option 82 information (when the system is not configured to do DHCP Relay), see Subscriber Identification Using Option 82 Field);
– to forward DHCP requests to a RADIUS server first, and not send them to the DHCP server unless the RADIUS server confirms positive identification.

For these applications, DHCP snooping must be enabled on the SAP towards the subscriber;

2. From the DHCP server (ACK messages):
   – to remove the Option 82 field toward the client
   – to build a dynamic DHCP lease state table for security purposes, see section DHCP Lease State Table
   – to perform Enhanced Subscriber Management, refer to Triple Play Enhances Subscriber Management

   For these applications, DHCP snooping must be enabled on both the SAP or SDP towards the network and the SAP towards the subscriber.

   A major application for DHCP response snooping in the context of Triple Play is security: A malicious user A could send an IP packet (for example, requesting a big video stream) with as source the IP address of user B. Any return packets would be sent to B, and thus potentially jam the connection to B.

   As the snooped information is coming straight from the operator's DHCP server, it is considered reliable. The BSA and BSR can use the snooped information to build anti-spoofing filters, populate the ARP table, send ARP replies, and so on.

3.2.5 DHCP Lease State Table

The DHCP lease state table has a central role in the BSA operation (Figure 23). For each SAP on each service it maintains the identities of the hosts that are allowed network access.

Figure 23 **DHCP Lease State Table**

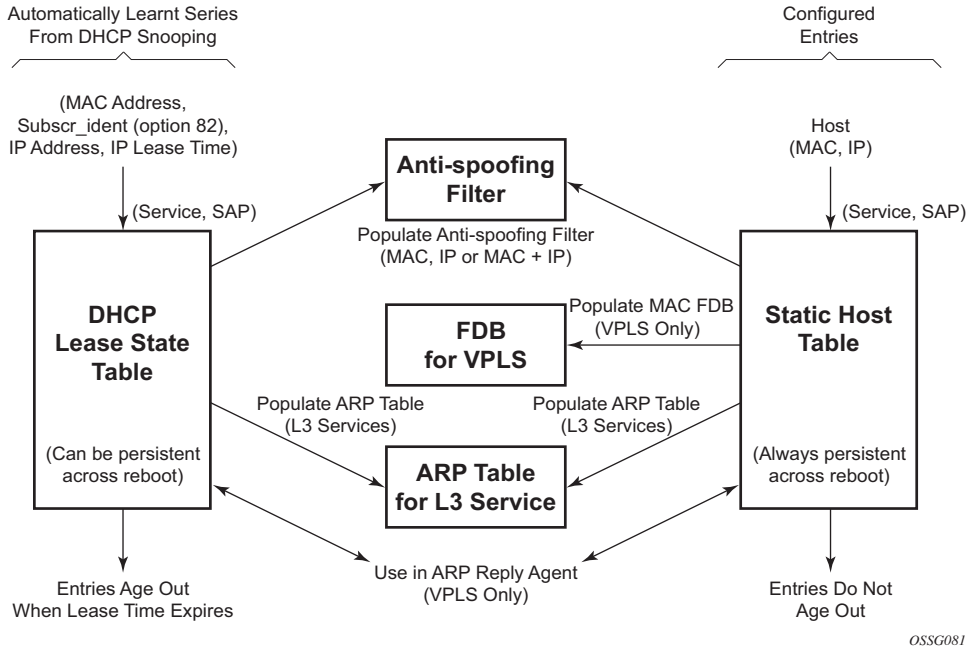

When the command **lease-populate** is enabled on a SAP, the DHCP lease state table is populated by snooping DHCP ACK messages on that SAP, as described in the DHCP Snooping section.

Entries in the DHCP lease state table remain valid for the duration of the IP address lease. When a lease is renewed, the expiry time is updated. If the lease expires and is not renewed, the entry is removed from the DHCP lease state table.

For VPLS, DHCP snooping must be explicitly enabled (using the **snoop** command) on the SAP and/or SDP where DHCP messages requiring snooping ingress the VPLS instance. For IES interfaces and VPRN IP interfaces (VPRN is supported on the 7750 SR only), using the **lease-populate** command also enables DHCP snooping for the subnets defined under the IP interface. Lease state information is extracted from snooped or relayed DHCP ACK messages to populate DHCP lease state table entries for the SAP or IP interface.

For IES and VPRN services, if ARP populate is configured, no statics ARPs are allowed. For IES and VPRN services, if ARP populate is not configured, then statics ARPs are allowed.

The retained DHCP lease state information representing dynamic hosts can be used in a variety of ways:

- To populate a SAP based anti-spoof filter table to provide dynamic anti-spoof filtering — Anti-spoof filtering is only available on VPLS SAPs, or IES IP or VPRN IP interfaces terminated on a SAP.
- To populate a VPLS SAP-based arp-reply-agent table to provide dynamic ARP replies using the dynamic hosts IP assigned IP address and learned MAC address — ARP reply agent functionality is only available for static and dynamic hosts associated with a VPLS SAP. arp-reply-agent is supported on the 7450 ESS only.
- To populate the system's ARP cache using the **arp-populate** feature — **arp-populate** functionality is only available for static and dynamic hosts associated with IES and VPRN SAP IP interfaces.
- To populate managed entries into a VPLS forwarding database — When a dynamic host's MAC address is placed in the DHCP lease state table, it automatically populates into the VPLS forwarding database associated with the SAP on which the host is learned. The dynamic host MAC address overrides any static MAC entries using the same MAC and prevent learning of the MAC on another interface (implicit MAC pinning on the 7450 ESS). Existing static MAC entries with the same MAC address as the dynamic host are marked as inactive but not deleted. If all entries in the DHCP lease state associated with the MAC address are removed, the static MAC may be populated. New static MAC definitions for the VPLS instance can be created while a dynamic host exists associated with the static MAC address. To support the Routed CO model, see to Routed Central Office (CO).
- To support enhanced Subscriber management, see to RADIUS Authentication of Subscriber Sessions.

If the system is unable to execute any of these tasks, the DHCP ACK message is discarded without adding a new lease state entry or updating an existing lease state entry; and an alarm is generated.

3.2.6 DHCP and Layer 3 Aggregation

3.2.6.1 DHCPv4 Snooping

The default mode of operation for DHCP snooping is that the DHCP snooping agent instantiates a DHCP lease state based on information in the DHCP packet, the client IP address and the client hardware address.

The mode of operation can be changed for DHCP snooping so that the Layer 2 header MAC address is used instead of the client hardware address from the DHCP packet for the DHCP lease state instantiation. This mode is selected by enabling l2-header in the **lease-populate** configuration command at the DHCP level. Because SR OS routers do not have the ability to verify the DHCP information (both the **src-ip** and **src-mac** of the packet are those of the previous relay point) anti-spoofing must be performed at the access node prior to the SR OS routers. This mode provides compatibility with MAC concentrator devices, and cable modem termination system (CMTS) and WiiMAX Access Controller (WAC).

A configuration example of a cable/wireless network together with subscriber management is shown in Figure 24. The subnet used to connect to the CMTS/WAC must be defined as a subnet in the subscriber interface of the Layer 3 CO model under which the hosts is defined. This means that all subscriber lease states instantiated on BSR must be from a “local” subscriber-subnet, even if those are behind the router, as there is no additional Layer 3 route installed pointing to them.

The important items to notice are static hosts at the subscriber interface side:

- IP-only static host pointing to CMTS/WAC WAN link is needed to allow BSR to reply to ARP requests originated from CMTS/WAC.
- IP-MAC static host pointing to CMTS/WAC access-facing interface is required to provide BSR with an arp entry for the DHCP relay address.

When dual-homing is used the CMTS/WAC may be configured with the same MAC for both upstream interfaces. If that is not possible the BSR can be configured with an optional MAC address. The BSR then uses the configured MAC address when instantiating the DHCP lease states.

Figure 24 **CMTS/WAC Network Configuration Example**

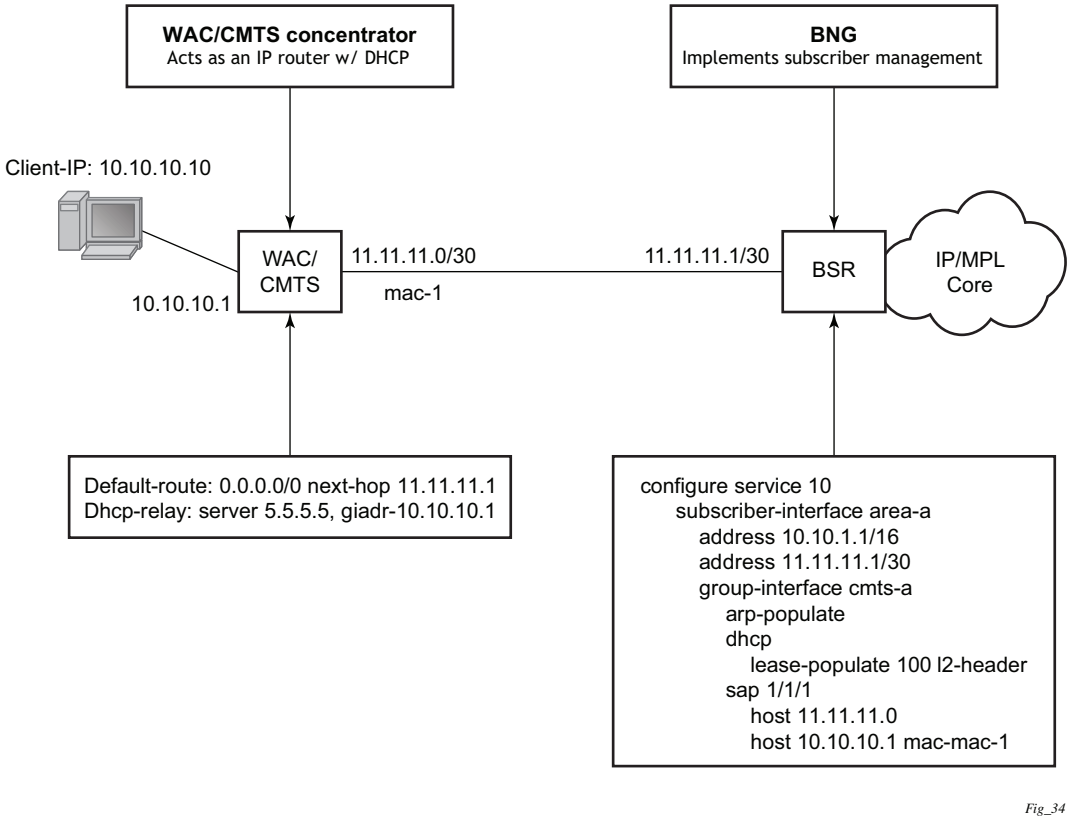

3.2.6.2 DHCPv6 Snooping

Similar to DHCPv4, the subscriber interface SAP must be on the data path between the subscriber host and the DHCPv6 server. The SAP snoops the DHCPv6 message exchanged between the server and the client. An ESM host is created upon snooping a “reply” message from the DHCPv6 server.

DHCPv6 messages differ from a DHCPv4 messages because it is not mandatory to have the client MAC inside the header or options. The DUID in the DHCPv6 option can be a random generated number instead of the subscriber host’s MAC. The source MAC of the DHCPv6 Ethernet header cannot be used either, because a Layer 3 aggregation network replaces the client’s MAC via routing. From the perspective of the BNG, all DHCPv6 message from the same downstream router has the same source MAC. By default, the BNG utilizes the DHCPv6 Ethernet header source MAC as a host entry identifier. Therefore, it is mandatory to use the interface ID in addition

to the source MAC to identify a host individually. If the interface ID is unavailable, it is possible to use python to copy another unique ID, such as DUID or remote ID, into the interface id. The interface ID option must be on the same level as the relay forward header. Together, the interface ID and the DHCP relay MAC address are used as an identifier internally in the BNG.

If the interface ID option is on the subscriber native DHCP message (such as solicit), it is simply ignored.

The downstream router must resolve the BNG MAC before it is able to route traffic to the BNG. Traditionally, a BNG sends router advertisement to directly connected hosts to help them resolve their default gateway and MAC address. However, routers differ from hosts and neighbor advertisements are used to resolve the neighbor's MAC instead. The downstream router has two options when programming the BNG as the next hop. It can either use the BNG subscriber interface link-local address or the subscriber interface GUA address. If an IPv6 prefix was configured on the BNG subscriber interface, then the downstream router must use the BNG link local as the next hop. If the subscriber interface is configured as an IPv6 address, then the downstream router can configure the GUA or the link-local as the next hop. In order to forward traffic bidirectionally, the downstream router interface must be modeled as a static host with both IP and MAC. IP-only static host is not supported.

The DHCP relay agent use one of its interface as a IP source address in the DHCP relay-forward message. The BNG forwards a DHCP relay-reply message from the DHCP server back to the relay agent using that exact same source IP address. There are restrictions for the IP source address used by the DHCP relay agent and it depends if the relay agent is a few hops way or is directly connected to the BNG. In the case where relay agent is a few hops away, the source address used by the relay agent must not fall under the subnet or prefix range configured on the subscriber interface. For example, the loopback or the egress interface address of the DHCP relay agent can be used instead. To forward the DHCP6 relay-forward message back to the relay agent, simply add a static route for the relay agent source IP address. The static route has the static IPv6 host as the next hop. In the case where the relay agent is directly connected to the BNG, there are two options. In the first option, the IPv6 static host configured on the BNG is actually an interface on the relay agent. If the relay agent use this as the relay-forward source address, no additional configuration is required on the BNG to forward relay-reply back to the relay-agent. The other option is to use an interface address on the relay agent which does not fall under the subnet or prefix under the subscriber interface. Similar to the scenario where the relay-agent is a few hops away, a static route is required to forward the DHCP relay-reply message back to the relay agent. Again the static route must use the IPv6 static host as the next hop.

A default host is supported for IPv6 host as well. It is generally used as a failover mechanism where the host can continue to forward traffic without a host entry on the backup BNG.

For IPv6 ESM host, it is mandatory that each host have a unique /64 prefix. Service providers who wish to share the /64 prefix among multiple WAN host can utilize the DHCPv6 filter **bypass-host-creation na** option. All bypass hosts in general require a default host creation as well.

While ESM hosts are subject to QoS and filters rules specified in sub-profile and sla-profile, default-host follows the QoS and filters specified directly on the subscriber SAP.

DHCP6 filters perform actions based on the options inside the relay-forward DHCP message. The options must be set on the innermost level, such as, DHCP solicit. The filter ignores those options set on relay-forward levels.

ESMv6 host created via DHCP snooping is not supported with the following:

- WLAN-GW
- DHCPv6 Proxy
- Wholesale or Retail
- SHCV
- L2-aware NAT
- GRT leaking or Extranet

3.2.7 Local DHCP Servers

3.2.7.1 Overview

The local-dhcp-server functions only if there is a relay (gateway) in front of it. Either a GI address is needed to find a subnet or Option 82, which is inserted by the relay, to perform authentication in the local-user-db.

A local DHCP server functions only if there is a relay agent (gateway) in front of it. The local DHCP server must be configured to assign addresses in one of the following ways:

1. Use a local user database authentication (**user-db** *local-user-db-name*)

   The host is matched against the specified local user database. A successful user lookup should return information on one of the following valid addresses:

   - A fixed IP address — The IP address should not overlap with the address ranges configured in the local dhcp server.
   - A pool name — A free address of any subnet in that pool is offered.

- use-gi-address [scope subnet | pool] — The GI address is used to find a matching subnet. When scope is subnet, an address is allocated in the same subnet as the GI address only. When scope is pool, an address is allocated from any subnet within a local pool once that pool has been selected based on matching the “giaddr” field in the DHCP message with any of the configured subnets in the pool.
- use-pool-from-client — The pool name specified in the DHCP client message options and added by the DHCP relay agent is used. A free address of any subnet in that pool is offered.

When no valid address information is returned from the local user database lookup, no IP address is offered to the client.

2. Without local user database authentication (**no user-db**).

   One or both address assignment options must be configured:

   - Use a pool name (use-pool-from-client)
     The pool name specified in the DHCP client message options and added by the DHCP relay agent is used. A free address of any subnet in that pool is offered.
   - Use the gi address (use-gi-address [**scope subnet** | **pool**])
     The gi address is used to find a matching subnet. When scope is subnet, an address is allocated in the same subnet as the GI address only. When scope is pool, an address is allocated from any subnet within a local pool once that pool has been selected based on matching the “giaddr” field in the DHCP message with any of the configured subnets in the pool.

When both options are configured and a pool name is specified in the DHCP client message options, then the **use-pool-from-client** option has precedence over the **use-gi-address** option.

Note: The local DHCP server will not allocate any address if none of the above options are configured (**no user-db**, **no use-gi-address**, **no use-pool-from-client**).

Options and identification strings can be defined on several levels. In principle, these options are copied into the DHCP reply, but if the same option is defined several times, the following precedence is taken:

1. user-db host options
2. Subnet options
3. Pool options
4. From the client DHCP request.

A local DHCP server must be bound to a specified interface by referencing the server from that interface. The DHCP server is then addressable by the IP address of that interface. A normal interface or a loopback interface can be used.

A DHCP client is defined by the MAC address and the circuit ID. This implies that for a certain combination of MAC and circuit ID, only one IP address can be returned. The same address is returned if a re-request is made.

Typically, the DHCP server can be configured to perform as follows:

- When a user-db is specified, a host lookup is performed in the local-user-db and the local user-db host defines how to get the IP address or pool to get the IP address from:
  - A fixed IP address.
  - A pool — A free address of any subnet in that pool is offered.
  - use-pool-from-client — The pool name is taken from Option 82 vendor-specific sub-option. (If not present, proceed to the next bullet.)
  - use-gi-address — The gi-address is used to find a matching subnet.
- If no IP address, subnet, or pool is found (see the above points), no IP address is offered to the client. If a subnet is found, an available address from the subnet is offered to the client. If a pool is found, an available address from any subnet of the pool is offered to the client.

Note: When the no use-gi, no use-pool-from-client and no user-db commands are specified, the DHCP server does not act.

3.2.7.2 Local DHCP Server Support

Local DHCP servers provide a standards-based full DHCP server implementation which allows a service provider the option of decentralizing the IP address management into the network. Local DHCP servers are supported on 7750 SR-7, and 7750 SR-12 models. The 7450 ESS routers use DHCP relay and proxy DHCP server functionality.

Three applications are targeted for the local DHCP server:

- Subscriber aggregation in either a single node (routed CO) or TPSDA.
- Business services — A server can be defined in a VPRN service and associated with different interfaces. Locally attached hosts can get an address directly from the servers. Routed hosts receives addresses through a relay point in the customer's network. This option is supported for the 7750 SR only.

- PPP clients — Either in a single node or a separate PPPoE server node and a second DHCP server node. The PPPoE server node may be configured to query the DHCP server node for an address and options to provide to the PPPoE client. The PPPoE server provides the information in PPP format to the client.

DHCP server scenarios include:

- DHCP servers can be integrated with Enhanced Subscriber Management (ESM) for DHCP clients, DHCP relays and PPPoE clients.
- A stand-alone DHCP server can support DHCP clients and DHCP relays.
- IPv4 is supported. DHCP servers provide increased management over the IPv4 address space across its subscriber base, with support for public and private addressing in the same router, including overlapped private addressing in the form of VPRNs in the same SR-series router.

DHCP servers are configurable and can be used in the bridged CO, routed CO, VPRN wholesaler, and dual-homing models.

3.2.8 DHCPv6

In the stateful auto-configuration model, hosts obtain interface addresses and configuration information and parameters from a server. Servers maintain a database that keeps track of which addresses have been assigned to which hosts. The stateful auto-configuration protocol allows hosts to obtain addresses, other configuration information or both from a server.

3.2.8.1 DHCPv6 Relay Agent

When the server unicast option is present in a DHCP message from the server, the option is stripped from the message before sending to the clients.

3.2.8.2 DHCPv6 Prefix Options

The prefix delegation options described in RFC 3633, *IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6*, provide a mechanism for automated delegation of IPv6 prefixes using DHCP. This mechanism is intended for delegating a long-lived prefix from a delegating router to a requesting router, across an administrative boundary, where the delegating router does not require knowledge about the topology of the links in the network to which the prefixes are assigned. For example, the delegating router can get a /48 prefix by DHCPv6 and assign /64 prefixes to multiple requesting routers. Prefix delegation is supported for the delegating router (not the requesting router).

3.2.8.3 Neighbor Resolution With DHCPv6 Relay

This is similar to ARP populate for IPv4. When it is turned on, an SR OS needs to populate the link-layer address to IPv6 address mapping into the neighbor database based on the DHCPv6 lease information received.

If the IPv6 address of the host doesn’t belong to the subnets of the interface, such a neighbor record should not be created. This could happen when there is a downstream DHCPv6 relay router or prefix delegation requesting router.

3.2.8.4 DHCPv6 Lease Persistency

DHCPv6 lease persistency is supported.

The following features are enabled:

- DHCPv6 lease information is reconciled to the standby.
- DHCPv6 lease information can be stored on a flash card.
- When rebooted, DHCPv6 lease information stored on a flash card can be used to re-populate the DHCPv6 table as well as the neighbor database if neighbor-resolution is enabled.

3.2.8.5 Local Proxy Neighbor Discovery

Local proxy neighbor discovery is similar to local proxy ARP. It is useful in the residential bridging environment where end users are not allowed to talk to each other directly.

When local proxy ND is turned on for an interface, the router:

- Responds to all neighbor solicitation messages received on the interface for IPv6 addresses in the subnet(s) unless disallowed by policy.
- Forwards traffic between hosts in the subnet(s) of the interface.
- Drops traffic between hosts if the link-layer address information for the IPv6 destination has not been learned.

3.2.8.6 IPv6oE Hosts Behind Bridged CPEs

This feature adds support for dual-stack IPoE hosts behind bridged clients, receiving globally-routable address using SLAAC or DHCPv6. The feature also provides configurable support for handing out /128 addresses to bridged hosts from same /64 prefix or a unique /64 prefix per host. Bridged hosts that share the same /64 prefix are required to be all SLAAC hosts or DHCPv6 hosts, and are required to be associated with the same SAP. For SLAAC based assignment, downstream neighbor-discovery is automatically enabled to resolve the assigned address.

3.2.8.7 IPv6 Link-Address Based Pool Selection

This feature provides the capability to select prefix pools for WAN or PD allocations based on configured link addresses. The scope of selection is the pool or a prefix range within the pool.

3.2.8.8 IPv6 Address and Prefix Stickiness

This feature extends lease identification criterion beyond DUID (default) for DHCPv6 leases held in the lease database for a configured period after the lease times out. DHCPv6 leases can be held in the lease database for a configurable period of time, after the lease time has expired. A large configured timeout value allows for address and prefix “stickiness”. When a subscriber requests a lease via DHCPv6 (IA_NA or PD), existing lease is looked-up and handed out. The lease identification match criterion has been extended beyond DUID to also include interface-id or interface-id and link-local address.

3.2.8.9 IPv4/v6 Linkage for Dual-Stack Hosts or Layer 3 RGs

In case of dual-stack Layer 3 RGs or dual-stack hosts behind Layer 2 CPEs, it is beneficial to have the ability to optionally link Ipv6oE host management to DHCP state for v4. This feature provides configurable support to use circuit-id in DHCPv4 option-82 to authenticate DHCPv6. Also, a SLAAC host is created based on DHCPv4 authentication if RADIUS returns IPv6 framed-prefix. IPv6oE host is deleted when the linked IPv4oE host is deleted. In addition, with v4 and v6 linkage configured, sending of unsolicited unicast RA towards the client can be configured when v4 host state is created and IPv6 is configured for the client. The linkage between IPv4 and IPv6 is based on SAP and MAC address. The sharing of circuit-id from DHCPv4 for authentication of DHCPv6 (or SLAAC) allows the SR series router to work around lack of support for LDRA on Access-nodes.

3.2.8.10 Host Connectivity Checks for IPv6

This feature provides support to perform SHCV checks on the global unicast address (assigned via SLAAC or DHCPv6 IA_NA) and link-local address of a L3 RG or a bridged host. SHCV uses IPv6 NS and NA messages. The configuration is similar to IPv4 support in SHCV. The **host-connectivity-check** command is extended to be configured for IPv6 or both IPv4 and IPv6.

3.2.9 Lease Query

A lease query by client-id as defined in RFC 5007, *DHCPv6 Lease Query*, is supported for Local DHCPv6 servers. For security reasons this must be explicitly enabled via the CLI command **allow-lease-query**. The user identification option must be set to DUID (default value) for lease-query to work. Lease query by address is not supported. It is not possible to filter out leases with the link address, the server always returns all addresses for a client. The Relay Data and Client Link options are not supported and are not returned.

3.2.10 DHCPv6 to Server Option

A DHCPv6-relay **user-db** can be configured on an IES or VPRN IP interface. This allows the insertion of a **to-client** and **to-server** option on the client's DHCPv6 message. The VPRN or IES IP interface must be the first relay agent for the subscriber. The interface must also have **lease-populate** enabled. The interface can store up to eight leases per DHCPv6 (relay) reply. If the interface contains both RADIUS and user-db, RADIUS always takes precedence. The **to-client** and **to-server** option is inserted in the inner most level of the DHCPv6 packet.

ESM IPoE subscriber can also utilize the **user-db** to insert **to-server** option. IPoE ESM subscribers are not limited to only **to-client** and **to-server** option, and can utilize all other parameters configurable under **config>subscr-mgmt>loc-user-db>ipoe>host>** in LUDB.

3.2.11 Flexible Host Identification in LUDB Based on DHCPv4/v6 Options

Host identification plays a critical role during the assignment of the parameters to the host via LUDB. The parameters that can be assigned to the subscriber host can range from the IP addressing parameters and the subscriber identification string all the way to the parameters that define the service to which the subscriber is entitled.

LUDB access in the context of IPoE hosts is triggered by DHCP messages passing through the interface on which the LUDB access is configured. This is true regardless of the direction of the DHCP message flow (ingress/egress).

The parameters that define the characteristics of the host are represented by an LUDB host *entry*. The parameters in the LUDB entry can be unique for each individual host, or they can be shared for a group of hosts. In the former case, the identification field for the LUDB host entry must be host specific while in the latter case the identification field for LUDB host entry could be derived from DHCP options that are common to a set of host.

The host identification in the LUDB can be based on a fixed set of predefined fields within the 7750 SR and 7450 ESS. If this predefined set of fields is not flexible enough, a custom identification field can be constructed from the DHCP options that are processed by the Python script. Once this custom identifier is constructed, its value can be preserved for the duration of the DHCP transaction and it is used by the LUDB for the host identification.

An example of how this can be used is the following:

A Python script is installed in 7750 SR and 7450 ESS. This Python script intercepts incoming DHCP messages on the access side (Discover/Solicit/Request/Renew/Rebind) and consequently creates a host identification string based on DHCP options in the packet. This string then is cached and used for host identification in LUDB in both directions (access ingress and network ingress).

This functionality is supported for DHCPv4/DHCPv6 hosts.

3.2.12 DHCP Caching

Subscriber host **identification** via LUDB is performed upon the arrival of the incoming DHCP messages on both, the access and the network side, while the host **instantiation and ESM string assignment** is performed only during the processing of the DHCP ACK/Reply messages. In other words, if Python without the caching is used for subscriber host identification and classification (into proper service class by means of deriving ESM strings), the DHCP options required for host identification must be present in all DHCP messages, even the ones sent by the DHCP servers. However, DHCP servers are not required to echo DHCP options sent by the clients and relay-agents. Consequently, the missing options from the server side would cause the subscriber host instantiation to fail.

To remedy this situation and cover all deployments models (even the ones where the DHCP options are not echoed back by the DHCP servers), a **caching** mechanism is introduced whereby the results of the Python processing on ingress access are locally stored in 7750 SR and 7450 ESS. This ensures that the information about the subscriber host is readily available when the DHCP packet from the DHCP server arrives. Furthermore, since we already have the cached information, no additional Python processing on the network ingress is needed.

The caching is performed in a *DHCP Transaction Cache (DTC)*, which is accessible to Python and to the ESM module. Python writes the result of its processing to it and the Enhanced Subscriber Management (ESM) module within 7750 SR and 7450 ESS can access those results.

The cache entries are relatively short lived, with the lifetime of a DHCP transaction. DHCP transaction is defined as a pair of DHCP messages that have the same DHCP transaction ID number (<Discover, Offer>, <Request, Ack>, <Solicit, Advertise>, <Request, Reply>, <Renew, Ack>, and so on).

3.2.13 Flexible Creation of DHCPv4/6 Host Parameters

One of the facilities for flexible creation and assignment of subscriber host parameters is through Python scripting.

There are two models that allow assignment of the subscriber hosts parameters based on the Python processing, one without the utilizing the internal cache (DTC) and the other with the internal cache (DTC).

- Without utilizing the DTC, Python can process options in DHCP ACK message, derive the subscriber host parameters based on those options and consequently insert those parameters in a pre-configured DHCP option (defined in *sub-ident-policy*). The ESM module can be then instructed to extract those parameters and consequently instantiate the host with proper service levels. The drawback of this solution is that the DHCP server may not return all DHCP (v4 and v6) options that clients and relay-agents originally transmitted. Since those options are needed for subscriber parameter determination (but may be absent in the DHCP ACK/REPLY messages when the Python script is run), this solution falls short of covering all deployment cases. In addition, the range of parameters that can be assigned to a subscriber host in this fashion is smaller than the set of parameters utilizing the DTC.
- The internal cache (DTC) allows us to store the result of Python processing. The result is stored during the lifetime of the DHCP transaction. This method of string assignment does not rely on the DHCP server ability to return client's options, DHCPv4 and DHCPv6.

Parameters (ESM strings, IP addresses, and so on) present in the DTC have priority over any other source that is providing overlapping parameters when it comes to ESM processing. In other words, if the same parameter is provided via DTC (Python), LUDB and RADIUS, the one provided via DTC is in effect. This prioritization occur automatically without any additional CLI.

For example, if the IPv4 address is provided via DTC during DHCP Discovery processing, then the mode of operation for this host is proxy-to-dhcp (ESM terminates DHCP, without going to the server), regardless of whether the IP address is also provided via LUDB or RADIUS.

This functionality is supported for DHCPv4/DHCPv6 hosts.

3.2.14 Python DTC Variables and API

The following are the Python variables and APIs related to DTC:

Subscriber Host Identification

alc.dtc.derivedId — A read or write (from the Python perspective) string to store the LUDB lookup key for subscriber host identification. This key is derived from the contents of the packet. This string is used as a *match* criteria in LUDB. The derived-id can only be used when the lookup is performed in ESM. If the LUDB is attached to the local DHCP server, then the lookup based on the derived-id cannot be performed as the DHCP server has no means to derive such an ID from the DHCP message.

Caching Any Data During the Lifetime of a Transaction

alc.dtc.store(key,value) — The operator can store any data he desires in one or more entries. The key can be any arbitrary string (printable ASCII characters), up to 32 bytes. The value part is ‘unlimited’ (memory permitting) in size.

alc.dtc.retrieve(key) — Retrieve data from the DTC. The key must be an existing key, which is a string consisting of printable ASCII characters, up to 32 bytes.

For example, this can be used to cache the DHCP options that the client inserts but the server does not echo back. Those options can still be retrieved in 7750 SR and 7450 ESS via cache in case that their presence is needed for any reason.

The lifespan of the cached data is tied to a DHCP transaction (a pair or corresponding DHCP messages flowing in opposite direction).

ESM Related Parameters (ESM strings, routing context)

DTC provides an API to supply a subset of configuration parameters that can otherwise come from RADIUS and LUDB and are used by the ESM code to setup the subscriber host.

DTC parameters as defined below should not be considered as DHCP options that can be blindly returned to the DHCP client, but instead they should be considered as real configuration settings. For example, the lease-time option is used in LUDB to enforce the lease time for the client. As such, the ESM keeps state of the lease-time. The following parameters can be used to setup a subscriber host:

alc.dtc.setESM (key-from-below, value) => store data that is used by ESM. This data is write-only.

The keys are predefined (**only** these can be used) and are described in Table 38. These keys are read-only static variables.

The LUDB column indicates the configuration option under the **config>subscr-mgmt>loc-user-db>ipoe>host** context in LUDB.

Table 38 **ESM-Related Python Variables**

DTC Variable	Type	LUDB	RADIUS Attribute	Comment
alc.dtc.subIdent	string	identification-strings >subscriber-id	Alc-Subsc-ID-Str	
alc.dtc.subProfileString	string	identification-strings >sub-profile-string	Alc-Subsc-Prof-Str	
alc.dtc.slaProfileString	string	identification-strings >sla-profile-string	Alc-SLA-Prof-Str	
alc.dtc.spiSharingGroupId	integer	identification-strings >spi-sharing-group-id	Alc-SPI-Sharing-Id	
alc.dtc.ancpString	string	identification-strings >ancp-string	Alc-ANCP-Str	
alc.dtc.appProfileString	string	identification-strings >app-profile-string	Alc-App-Prof-Str	
alc.dtc.intDestId	string	identification-strings >inter-dest-id	Alc-Int-Dest-Id-Str	
alc.dtc.catMapString	string	identification-strings >category-map-name	Alc-Credit-Control-CategoryMap	
alc.dtc.ipAddress	string	address	Framed-IPAddress	
alc.dtc.dhcp4DefaultGateway	string	options>default-router	Alc-Default-Router	
alc.dtc.subnetMask	string	address	Framed-IPNetmask	
alc.dtc.ipv4LeaseTime	integer	options>lease-time	Alc-Lease-Time	
alc.dtc.ipv4PrimDns	string	options>dns-server	Alc-Primary-Dns Client-DNS-Pri	
alc.dtc.ipv4SecDns	string	NA	Alc-Secondary-Dns Client-DNS-Sec	
alc.dtc.primNbns	string	options>netbios-name-server	Alc-Primary-Nbns RB-Client-NBNSPri	
alc.dtc.secNbns	string	NA	Alc-Secondary-Nbns RB-Client-NBNSSec	
alc.dtc.msapGroupInterface	string	msap-defaults>group-interface	Alc-MSAP-Interface	

Table 38 **ESM-Related Python Variables (Continued)**

DTC Variable	Type	LUDB	RADIUS Attribute	Comment
alc.dtc.msapPolicy	string, integer	msap-defaults>policy	Alc-MSAP-Policy	
alc.dtc.msapServiceId	string, integer	msap-defaults>service	Alc-MSAP-Serv-Id	
alc.dtc.retailServiceId	string	Retail-service-id	Alc-Retail-Serv-Id	
alc.dtc.ipv6Address	string	ipv6-address	Alc-Ipv6-Address	
alc.dtc.ipv6DelegatedPrefix	string	ipv6-delegated-prefix	Delegated-IPv6-Prefix	
alc.dtc.ipv6SlaacPrefix	string	ipv6-slaac-prefix	Framed-IPv6-Prefix	
alc.dtc.ipv6WanPool	string	ipv6-slaac-prefix-pool	Framed-IPv6-Pool	
alc.dtc.ipv6PrefixPool	string	ipv6-delegated-prefix-pool	Alc-Delegated-IPv6-Pool	
alc.dtc.ipv6DelegatedPrefixLength	integer	ipv6-delegated-prefix-len	Alc-Delegated-IPv6-Prefix-Length	
alc.dtc.accountingPolicy	string	acct-policy	N/A	
alc.dtc.dhcpv4GIAddr	string	gi-address	N/A	
alc.dtc.dhcv4ServerAddress	string	server	N/A	
alc.dtc.dhcp4SrcAddr	string	N/A	N/A	
alc.dtc.dhcp4Pool	string	address>pool	Framed-Pool Ip-Address-Pool-Name	prim | sec (“|” – delimiter)
alc.dtc.linkAddress	string	link-address	N/A	
alc.dtc.dhcp6SrcAddr	string	N/A	N/A	
alc.dtc.dhcv6ServerAddr	string	server6	N/A	
alc.dtc.setDhcpv6LinkAddr	string	link-address	N/A	This API applies only to regular numbered or unnumbered IPv6 interfaces (no ESM)

Table 38 **ESM-Related Python Variables (Continued)**

DTC Variable	Type	LUDB	RADIUS Attribute	Comment
alc.dtc.setDhcpv6ServerAddr	string	server6	N/A	This API applies only to regular numbered or unnumbered IPv6 interfaces (no ESM)
alc.dtc.ipv6PrimDns	string	options6>dns-server	Alc-Ipv6-Primary-Dns	
alc.dtc.ipv6SecDns	string	NA	Alc-Ipv6-Secondary-Dns	
alc.dtc.dhcpv6PreferredLifetime	integer	ipv6-lease-times>preferred-lifetime	Alc-v6-Preferred-Lifetime	
alc.dtc.dhcpv6RebindTimer	integer	ipv6-lease-times>rebind-timer	Alc-Dhcp6-Rebind-Time	
alc.dtc.dhcpv6RenewTimer	integer	ipv6-lease-times>renew-timer	Alc-Dhcp6-Renew-Time	
alc.dtc.dhcpv6ValidLifetime	integer	ipv6-lease-times>valid-lifetime	Alc-v6-Valid-Lifetime	

For example, an IP address is assigned to a DTC variable as a string:

alc.dtc.ipAddress = 192.168.0.10 — This is performed through the following ALU API: alc.dtc.setESM(alc.dtc.ipAddress**,** 192.168.0.10). The DTC logic then parses this variable and converts it into appropriate format for consumption by ESM code.

The values defined above are the ones that are mostly defined in the LUDB. Main use, however, is assigning ESM strings for the subscriber host instantiation phase during the processing of DHCP ACK/Reply messages. Consequently the Python script needs to be run only on DHCP Request messages (no need to run it on Discoveries for ESM string assignment, unless the LUDB derived ID is also needed).

DHCP options that are blindly returned to the DHCP client without the ESM code being aware of them cannot be configured via DTC. Such options should be configured via RADIUS (Alc-ToCLient-Dhcp-Options – IPv4 only) or they can be inserted directly into DHCP messages via Python (bypassing DTC).

Other possible uses for DTC variables are:

- Assigning routing context information via Python (service-id, msap, msap-policy, retail service-id, and so on). For example AN can insert certain hints in various DHCP options that would suggest to us (via Python) in which service context to place the subscriber-host.
- IP address assignment via Python (DTC). This would address the DTC-to-DHCP-Proxy case where Python script is invoked on DHCP Discovery/Solicit. For example:
  - Discover arrives
  - Python generates an IP address, for example based on some DHCP options
  - The script stores the IP address by using alc.dtc.setEsm(alc.dtc.ipAddress, 10.0.1.2);
  - After the script is finished, ESM starts processing the packet (no LUDB/RADIUS authentication configured)
  - ESM finds the IP address already in DTC and decides to handle all DHCP and execute proxy function instead of relay
  - ESM sends an offer with the address that python generated
  - DHCP options should be provided as well in this case (lease-times, and so on)
  - The same applies to the DHCP Request

3.2.14.1 DTC Debugging Facility

DTC debugging is part of the generic DHCP debugging facility that is enabled by the flowing commands:

debug router *router-id* **ip dhcp** — Enables DHCP debug on Layer 3 interfaces, including subscriber-interfaces

debug service id *service-id* **dhcp** — Enables DHCP debugging on capture SAP

If the DTC cache is populated via Python, the corresponding DTC entries are shown as part of the matching DHCP message debug.

3.2.15 Virtual Subnet for DHCPv4 Hosts

The **virtual-subnet** command in the **sub-if>dhcp** context allows the system to snoop and record the default router address in DHCP ACK messages for a DHCPv4 ESM host. The system can answer ping or traceroute request even if the default router address is not configured on the subscriber-interface.

This feature eliminates the need to configure every default-gw address on subscriber interface. Beside default router address, the system also calculates host's subnet by using an assigned address and the subnet mask option in ACK. Both recorded default router address and the subnet can be displayed with the **show service** *id* **virtual-subnet** command.

Every ESM subscriber only has one set of default router address and subnet.

3.2.16 Address Reservation for Sticky Leases

Address reservation for sticky leases adds support in local DHCP servers to provide IP address reservation for the assignment of sticky leases. These leases are pre-provisioned in the server (via SNMP or CLI) with a given set of user-identification parameters. This set of parameters must be unique for each pool to avoid duplicate leases. For management purposes, these leases also have an additional *hostname* parameter to easily retrieve them via SNMP. When a sticky lease is created, the corresponding requested IP address is allocated from the specified pool and this address afterwards can only be used by DHCP if they match the specified user identification parameters. It is not possible to make an existing DHCP lease sticky.

Sticky leases are persistent but not synchronized in multi-chassis synchronization. To support multi-chassis redundancy, a management system can allocate the lease on one 7750 SR, immediately retrieve the lease and populate it again on the redundant router. For this to work sticky leases should not be combined with any other allocation method (for example, regular DHCP leases, or Local Address Assignment because these methods can already allocate the address on the standby node).

This feature targets two scenarios:

- Simplified IP address reservation. It is not necessary to provision LUDB entries and exclude ranges for sticky leases. Once the lease is reserved in the local DHCP server (triggered by an external system via SNMP or CLI), the external system can then subsequently assign the corresponding IP address to a DHCP client that matches the configured lease. A use case of this related to virtual residential gateway (vRGW) application and is described in the vRGW section.

- Pool management without DHCP. In pure management cases, this provides an easy method to perform pool management without implementing DHCP-specific configurations. For example, a management platform can allocate an IP address from a pool via the sticky lease mechanism and then assign this to a static host without risking overlap.

3.3 Proxy DHCP Server

This section describes the implementation of proxy DHCP server capability to provide a standards-based DHCP server which will front-end to a downstream DHCP client, DHCP relay enabled devices, and interfaces with RADIUS to authenticate the IP host and subscriber, and obtains the IP configuration information for DHCP client devices.

The proxy DHCP server is located between an upstream DHCP server and downstream DHCP clients and relay agents when RADIUS is not used to provide client IP information.

Service providers can introduce DHCP into their networks without the need to change back-end subscriber management systems that are typically based around RADIUS (AAA). Service providers can support the use of DHCP servers and RADIUS AAA servers concurrently to provide IP information for subscriber IP devices (Figure 25).

Figure 25 **Typical DHCP Deployment Scenarios**

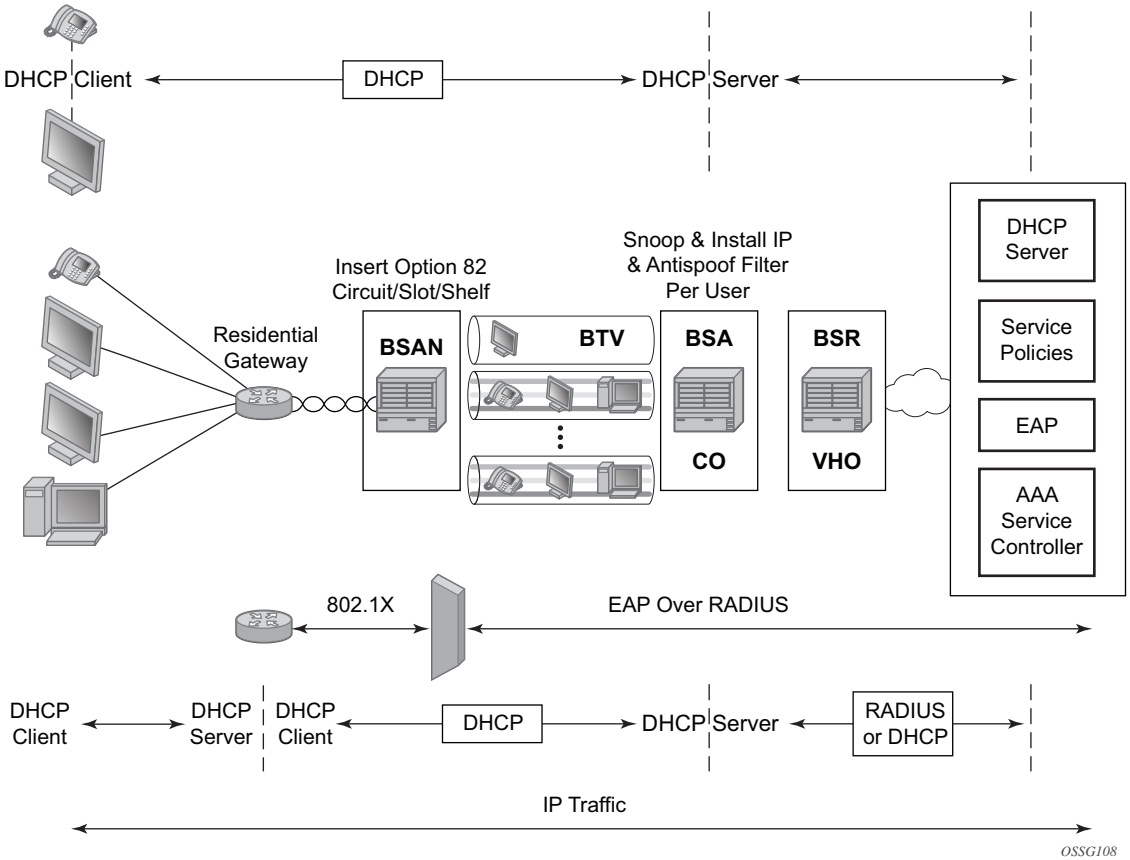

DHCP is the predominant client-to-server based protocol used to request IP addressing and necessary information to allow an IP host device to connect to the network.

By implementing DHCP, the complexity of manually configuring every IP device that requires connectivity to the network is avoided. IP devices with DHCP can dynamically request the appropriate IP information to enable network access.

DHCP defines three components that are implemented in a variety of device types:

- The DHCP client allows an IP device (host) to request IP addressing information from a DHCP server to enable access to IP based networks. This is typically found in:
  - End user notebooks, desktops and servers
  - Residential gateways and CPE routers
  - IP phones
  - Set-top boxes
  - Wireless access points
- The DHCP relay agent passes (relays) DHCP client messages to pre-configured DHCP servers where a DHCP server is not on the same subnet as the IP host. This feature optionally adds information into DHCP messages (Option 82) which is typically used for identifying attaching IP devices and their location as part of subscriber management. This is typically found in:
  - Residential gateways and CPE routers
  - DSLAMs
  - Edge aggregation routers
- The DHCP server receives DHCP client messages and is responsible for inspecting the information within the messages and determining what IP information if any is to be provided to a DHCP client to allow network access. This is typically found in:
  - Dedicated standalone servers
  - Residential gateways and CPE routers
  - Edge aggregation routers
  - Centralized management systems

DHCP is the predominant address management protocol in the enterprise community, however in the provider market PPP has traditionally been the means by which individual subscribers are identified, authenticated and provided IP addressing information.

The use of DHCP in the provider market is a growing trend for managing subscriber IP addressing, as well as supporting newer devices such as IP-enabled IP phones and set-top boxes. The majority of subscriber management systems rely heavily on RADIUS (RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*) as the means for identifying and authorizing individual subscribers (and devices), deciding whether they are allowed access to the network, and which policies should be put in place to control what the subscriber can do within network.

The proxy DHCP server capability enables the deployment of DHCP into a provider network, by acting as a proxy between the downstream DHCP devices and the upstream RADIUS based subscriber management system.

- Interact with downstream DHCP client devices and DHCP relay Agents in the path
- Interface with RADIUS to authenticate DHCP requests
- Receive all the necessary IP information to properly respond to a DHCP client
- Override the allocated IP address lease time, if necessary, for improved IP address management.

Figure 26 shows a typical DHCP initial boot up sequence with the addition of RADIUS authentication. The proxy DHCP server interfaces with downstream DHCP client devices and then authenticate upstream using RADIUS to a providers subscriber management system.

Figure 26 **Aggregation Network With DHCP to RADIUS Authentication**

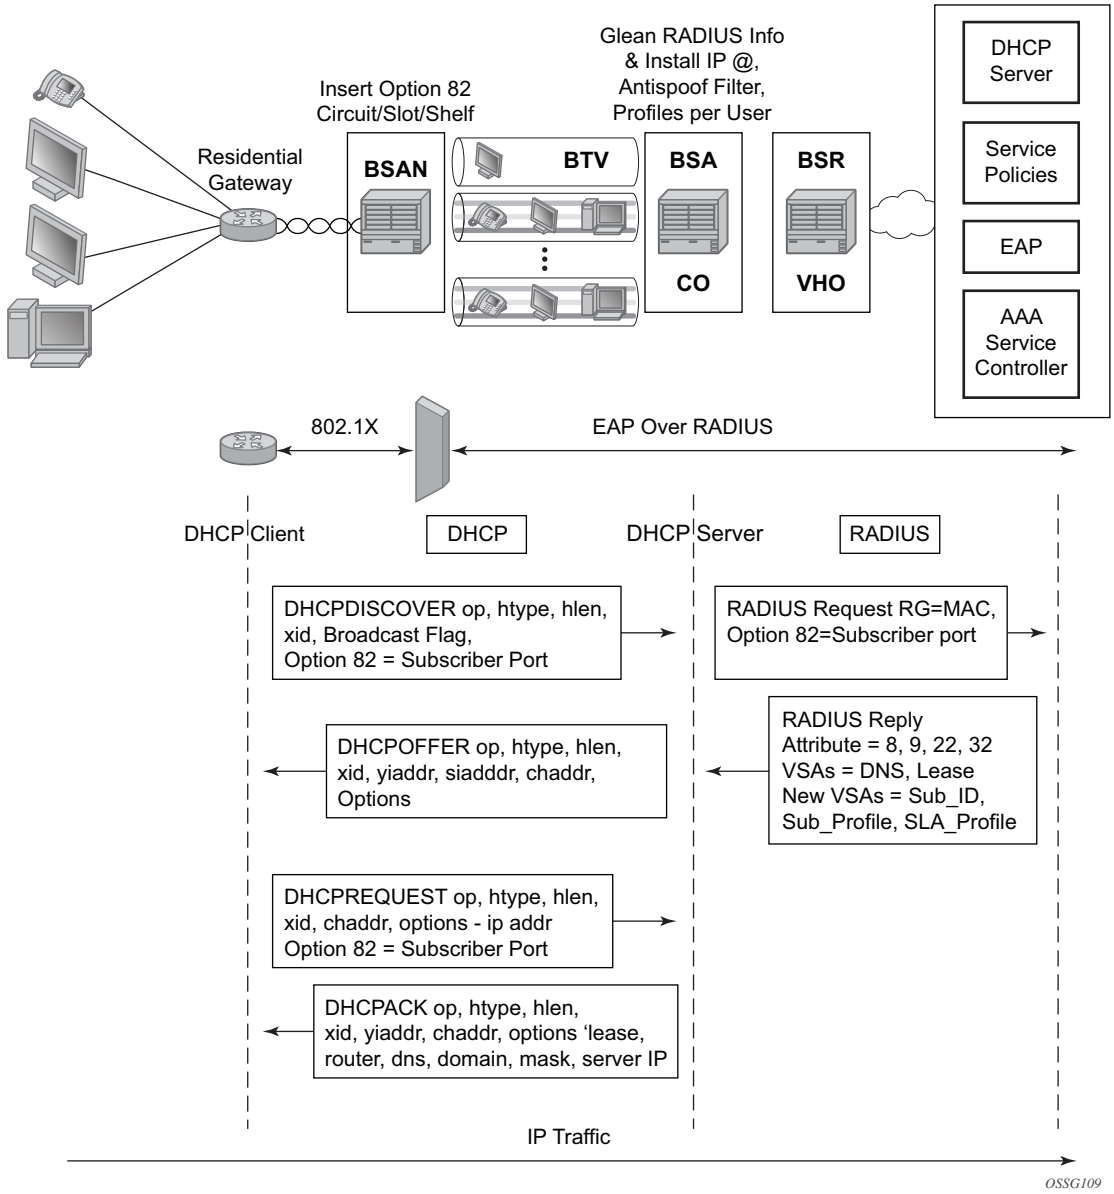

In addition to granting the authentication of DHCP hosts, the RADIUS server can include RADIUS attributes (standard and vendor-specific attributes (VSAs)) which are then used by the edge router to:

- Provision objects related to a given DHCP host — Subscriber and SLA policy
- Provide IP addressing information to a DHCP client
- Support the features that leverage DHCP lease state
  - Dynamic ARP population

  - ARP reply agent
  - Anti-spoofing filters
  - MAC pinning
- Leverage host-connectivity-verify to determine the state of a downstream IP host

This feature offers the ability for a customer to integrate DHCP to the subscriber while maintaining their existing subscriber management system typically based around RADIUS. This provides the opportunity to control shifts to an all DHCP environment or to deploy a mixed DHCP and RADIUS subscriber management system.

In order to maximize its applicability VSAs of legacy BRAS vendors can be accepted so that a network provider is not forced to reconfigure its RADIUS databases (or at least with minimal changes).

To receive data from the RADIUS server the following are supported:

- Juniper (vendor-id 4874) attributes 4 (Primary DNS server) and 5 (Secondary DNS server)
- Redback (vendor-id 2352) attributes 1 (Primary DNS) and 2 (Secondary DNS).
- Juniper attributes 6 and 7 (Primary and Secondary NetBIOS nameserver).
- Redback attributes 99 and 100 (Primary and Secondary NetBIOS nameserver).

The following attributes can be sent to RADIUS:

- Sending authentication requests: (from the DSL Forum) (vendor-id 3561), attributes 1 (Circuit ID) and 2 (Remote ID).
- DSL Forum attributes 129 and 130 (Actual Data Rate Upstream and Downstream), 131 and 132 (Minimum Data Rate Upstream and Downstream) and 144 (Access Loop Encapsulation).

The complete list of Nokia VSAs is available on a file included on the compact flash shipped with the image.

3.3.1 Local DHCP Servers

3.3.1.1 Terminology

- **Local 7750 SR and 7450 ESS DHCP Server** – Refers to the DHCP server instantiated on the local 7750 SR and 7450 ESS node.
- **Remote 7750 SR and 7450 ESS DHCP Server** – Refers to the DHCP server instantiated on the remote 7750 SR and 7450 ESS node (external to the local 7750 SR and 7450 ESS node).
- **3rd party DHCP server** – Refers to the DHCP server external to any 7750 SR and 7450 ESS node and implemented outside of 7750 SR and 7450 ESS.
- **Intercommunication link** – Refers to the logical link between dual-homed 7750 SR and 7450 ESS DHCP servers used for synchronizing DHCP lease states. Multi-chassis Synchronization (MCS) protocol runs over this link. Once this link is interrupted, synchronization of the leases between redundant DHCP servers is impaired. This link should be well protected with multiple underlying physical paths.
- **Local IP address-range and prefix** – Refers to the local failover mode in which the IP address-range and prefix is configured in dual-homed DHCP environment. The local keyword does not refer to the locality (local vs remote) of the server on which the IP address-range and prefix in configured, but rather refers to the ownership of the IP address-range and prefix. The DHCP server on which the local IP address-range and prefix is configured, owns this IP address-range and prefix and consequently is allowed to delegate the IP addresses and prefixes from it at any time, regardless of the state of the intercommunication link.
- **Remote IP address-range and prefix** – Refers to the remote failover mode in which the IP address-range and prefix is configured in dual-homed DHCP environment. The remote keyword does not refer to the locality of the server on which the IP address-range and prefix is configured, but rather refers to the ownership of the IP address-range and prefix. The DHCP server on which the remote IP address-range and prefix is configured, but does not own this IP address-range and prefix during normal operation and consequently is not allowed to delegate the IP addresses and prefixes from it. Only when the intercommunication link between the two nodes transition into particular (failed) state, the DHCP server is allowed to start delegating new IP addresses from the remote IP address-range and prefix.

- **IP address-range and prefix ownership** – Refers to 7750 SR and 7450 ESS DHCP servers that can delegate new leases from an IP address-range and prefix that it owns. For example, an IP address-range and prefix designated as remote is not owned by the DHCP server on which it is configured unless certain conditions are met. Those conditions are governed by the state of the intercommunication link.
- **IP address-range and prefix takeover** – Refers to 7750 SR and 7450 ESS DHCP servers that do not own an IP address-range and prefix can take over the ownership of this IP address-range and prefix under certain conditions. Once the ownership is taken, the new IP addresses can start being delegated from this IP address-range and prefix. Only the remote IP address-range and prefix can be taken over. Note that the takeover of an IP address-range and prefix has only local significance – in other words, the ownership is not taken away from some other DHCP server that has the same IP address-range and prefix designated as local. It only means that IP address-range and prefix that is configured as remote is available to takeover for new IP address delegation.
- **Local PPPoX Address Pools** – Refers to the method of accessing an IPv4/v6 address pool in 7750 SR and 7450 ESS DHCP4/6 server. For PPPoX clients, the IPv4/v6 addresses are allocated from those pools without the need for an intermediate DHCP relay-agent (7750 SR and 7450 ESS internal DHCP relay-agent). Although those pools are part of the local DHCP server in 7750 SR and 7450 ESS, the method of accessing them is substantially different than accessing local DHCP address pools for IPoE (DHCP) clients. IPoE (DHCP) and PPPoX hosts can share the same pool and yet each client type can access them in their own unique way:
  - IPoE client via DHCP messaging
  - PPPoE via internal API calls
- **Local PPPoX Pool Management** – Refers to the IPv4 address allocation/management for PPPoX clients independent of DHCP process (DHCP lease state). An IPv4 address allocated via local PPPoX Pool Management is tied to the PPPoX session. It is without the need for an internal DHCP relay-agent.

3.3.1.2 Overview

7750 SR and 7450 ESS DHCP server multi-homing ensures continuity of the IP address and prefix assignment and renewal processes when an entire 7750 SR and 7450 ESS DHCP server fails or in case of a failure of the active link that connects clients to one of the 7750 SR and 7450 ESS DHCP servers in the access part of the network. DHCP server multi-homing is an integral part of the overall subscriber management multi-chassis protection scheme.

DHCP server multi-homing can be implemented outside of the BNG, without subscriber management enabled. However, in the following text, it is assumed that the subscriber management multi-homing (SRRP, MC-LAG, subscriber synchronization) is deployed along with DHCP server multi-homing.

Although the subscriber synchronization process and the DHCP lease states synchronization process use the same synchronization infrastructure within 7750 SR and 7450 ESS (Multi Chassis Redundancy protocol), they are in essence two separate processes that are not aware of each other. As such, the mechanisms that drive their switchover are different. For example, the mechanism that drives subscriber switchover from one node to the other is driven by the access protection mechanism (SRRP/MC-LAG) while the switchover (or takeover) of the IP address-range and prefixes in a DHCP pool is driven by the state of the intercommunication link over which the leases are synchronized. The failure of an entire node makes those differences irrelevant because the access-link failure coincides with the intercommunication link failure and vice versa. However, link-only failures become critical when it comes to their interpretation by the protection mechanisms (SRRP, MC-LAG, DHCP server multi-homing). Regardless of nature of the failure, an overall DHCP server multi-chassis protection scheme must be devised in so that the two 7750 SR and 7450 ESS DHCP servers never allocate the same IP address and prefix to two different clients. Otherwise, IP address or prefix duplication ensues. Unique IP address and prefix allocation is achieved by making only one 7750 SR and 7450 ESS DHCP server responsible for IP address prefix delegation out of the shared IP address-range or prefix.

There are two basic models for DHCP server dual-homing:

- Shared IP address-ranges and prefixes are designated as local on one 7750 SR and 7450 ESS DHCP server and as remote on the other.

  In this case, the DHCP relays must point to both DHCP servers; the one configured with the local IP address-range and prefix as well as the one with the remote IP address-range and prefix.

  Under normal circumstances, the new IP addresses and prefixes can be only allocated from the DHCP server configured with the local IP address-range and prefix.

  The DHCP server configured with the remote IP address-range and prefix starts delegating new lease from it only when it declares that the redundant peer with the local IP address-range and prefix becomes unavailable.

  Detection of the peer unavailability is triggered by the failure of the intercommunication link which can be caused either by the nodal failure or simply by the loss of connectivity between the two nodes protecting each other. Thus, the loss of intercommunication link does not necessarily mean that the peering node is truly gone. It can simply mean that the two nodes became

isolated and unable to synchronize their DHCP leases between each other. In such environment, both nodes can potentially allocate the same IP address at the same time. To prevent this, additional intercommunication link states and associated timers are introduced to give the operator ample time to fix the problem.

For example, the DHCP server takes over the remote IP address-range and prefix once the MCLT period expires while the intercommunication link is in PARTNER-DOWN state. The PARTNER-DOWN state is entered after a preconfigured timer (partner-down-delay) expires. The consequence of these two additional timers (partner-down-delay and MCLT) is that the new IP address delegation from the remote (shared) IP address-range and prefix is not possible until the preset timers expire. This is needed and justified if the intercommunication link is interrupted, the nodes become isolated, and consequently, the DHCP lease state synchronization becomes impaired. On the other hand, if the DHCP server with local IP address-range and prefix becomes truly unavailable, those additional restoration times causes interruption in service since the new IP addresses from the remote IP address-range and prefix is not immediately available for delegation.

Only new IP address delegation from the remote IP address-range and prefix is affected by this behavior. The existing IP leases can be extended on both nodes at any time irrespective of whether the configured address-range and prefix is designated as local or remote.

To ensure uninterrupted service even for new lease delegation in this model (local-remote), two approaches can be adopted:

- Segment the IP address and prefix space so that each node has an IP address-range and prefix designated as local.
  For example, instead of designating IP address-range 10.10.10.0/24 as local on DHCP server A and as remote on DHCP server B, the 10.10.10./24 IP address is split into two: 10.10.10.0/25 and 10.10.10.128/25.
  The 10.10.10.0/25 would be designated as local on the DHCP server A and as remote on DHCP server B.
  The 10.10.10.128/25 would be designated as remote on the DHCP server A and as local on DHCP server B.
  In this fashion, one node is always available to assign new leases without any overlap.
- In case that only one shared IP address-rage/prefix is deployed, the operator can bypass the timers (partner-down-delay and MCLT) that are put in place in case that DHCP server nodes become isolated. This bypass of the timers can be achieved via configuration. In this case, a safe operation is warranted only if the operator is certain that the intercommunication link failure is caused by the nodal failure, and not the physical link failure between the two nodes.

• Shared IP address-range and prefix is designated as access-driven on both 7750 SR and 7450 ESS DHCP servers.

In this scenario, the shared IP address-range and prefix is owned by both nodes and the ownership is not driven by the state of the intercommunication link.

To avoid IP address duplication, only one DHCP server at any given time must be responsible for IP address assignment from this shared IP address-range and prefix.

This is ensured by the access protection mechanism (SRRP/MC-LAG) that provides a single active path from the clients to the one of the DHCP servers.

In case that clients have access to the same IP address-range and prefix on both DHCP servers at the same time, the IP address duplication may occur.

Consider the following case:

Two DHCP clients send DHCP Discovers in the following fashion:

- DHCP client A sends DHCP Discover to the DHCP server A
- DHCP client B sends DHCP Discover to the DHCP server B
- DHCP server A assigns IP address 10.10.10.10 to the DHCP client A
- DHCP server B assigns IP address 10.10.10.10 to the DHCP client B
- This is a legitimate scenario since the DHCP lease states are not synchronized until the DHCP lease assignment is completed.
- Just before the DHCP ACK is sent to the respective clients from both nodes, the DHCP lease sync messages are exchanged between the peers.
- DHCP servers do not wait for the reply to the sync message before they send the DHCP Ack to the client.
- Once the DHCP lease syn message is received from the peer, the DHCP server realizes that the IP lease already exist. In this case, the newer IP lease overrides the older.
- The result is that clients A and B use the same IP address and consequently the forwarding of the traffic is impaired.

In access-driven model, the ESM subscriber host must be collocated with the DHCP server. In other words, the DHCP server must be instantiated in redundant BNGs. The dhcp-relays must point to the respective local DHCP servers. There must be no cross-referencing of DHCP servers in this model. In addition, the IP address that the DHCP servers are associated with, must be the same on both DHCP servers. This is necessary to ensure uninterrupted service levels once the switchover in the access occurs.

For example:

- DHCP server A on BNG A is associated with the IP address 1.1.1.1 (for example loopback interface A on BNG A)
- DHCP server B on the peering BNG B is associated with the IP address 1.1.1.1 (configured in BNG B under loopback interface B)
- Dhcp-relay on BNG A points to the IP address 1.1.1.1—DHCP server in BNG A

– Dhcp-relay on BNG B points to the IP address 1.1.1.1—DHCP server in BNG B.

Consider the following when contemplating deployment of the two described models:

- Local-remote model is agnostic of the access protection mechanism. In fact, the access protection mechanism is not needed at all for safe operation.

  Fast takeover of the single shared (remote) IP address-range and prefix can be provided only in cases where the operator can guarantee that the intercommunication link failure is caused by the nodal failure (entire DHCP server node becomes unavailable). Fast takeover is provided by bypassing the partner-down-timer and MCLT.

  In case that multiple IP address-ranges and prefixes are deployed, bypass of the timers is not needed since the local IP address-ranges and prefixes are available on both nodes.
- Access-driven model allows a single IP address-range and prefix to be shared across the redundant DHCP server nodes. The access to the single DHCP server node from the client side is ensured by the protection mechanism deployed in the access part of the network (SRRP or MC-LAG).

3.3.1.3 DHCP Lease Synchronization

DHCP server leases are synchronized over Multi-Chassis Synchronization (MCS) protocol. A DHCP lease synchronization message is sent to the peering node just before the DHCP ACK/Reply is sent to the client.

For example, the message flow for DHCPv4 lease establishment is the following:

```
DHCP Discover        ; DHCP client — DHCP server
DHCP Offer           ; DHCP server — DHCP client
DHCP Request         ; DHCP client — DHCP server
DHCP Sync Message    ; DHCP server — via MCS to the peering DHCP server
DHCP Ack             ; DHCP server— DHCP client
```

DHCP server failover mechanism in the local-remote IP address-range and prefix model relies on the detection of the failure of the link over which DHCP states are synchronized (via MCS protocol). This link is normally disjoint from the access links towards the clients. MCS protocol normally runs over a direct link between the two redundant nodes (1) or over backbone links (2) in case that the direct link is not present. This is shown in Figure 27.

Figure 27 **Redundancy Model**

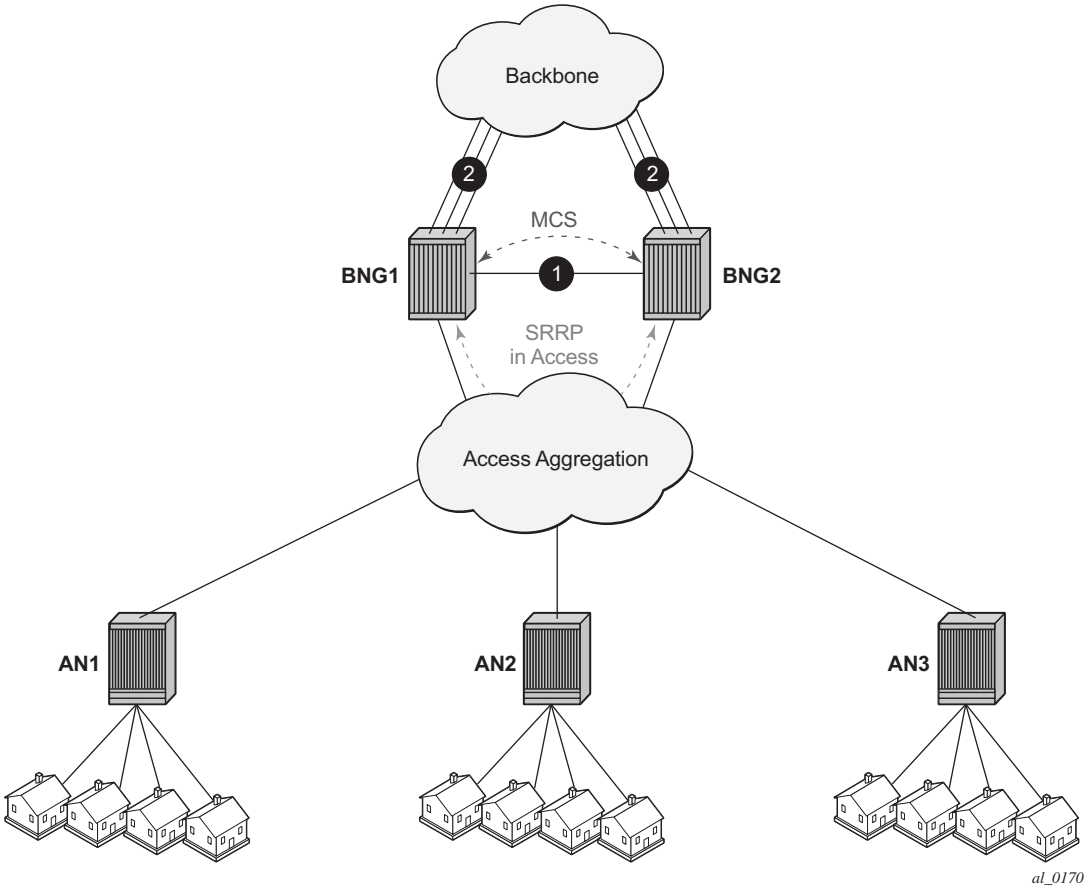

In the access-driven IP address-range and prefix model, the DHCP server address-ranges and prefixes are not tied to the state of the intercommunication link. Instead, the DHCP server selection for IP address assignment is only governed by the path selected by the path protection mechanism (SRRP/MC-LAG) deployed in the access part of the network.

3.3.1.4 Intercommunication Link Failure Detection

7750 SR and 7450 ESS DHCP server is a client of Multi-chassis Synchronization (MCS) application with 7750 SR and 7450 ESS. Once MCS transitions into out-of-sync state, the 7750 DHCP server redundancy assumes that there is a failure in the network. In essence, the DHCP server failure in dual-chassis configuration relies on the failure detection mechanism of MCS.

MCS runs over TCP, port 45067 and it is using either data traffic or keepalives to detect failure on the communication link between the two nodes. In the absence of any MCS data traffic for more than 0.5 seconds, MCS sends its own keepalive to the peer. If a reply is not received within 3 seconds, MCS declares its operational state as DOWN and the DB Sync state as out-of-sync. MCS consequently notifies its clients (DHCP Server being one of them) of this condition.

In essence, it can take up to 3 seconds before the DHCP client realizes that the interclasses communication link failed.

MCS Clients (applications) can optimally send their own proprietary keepalive messages to its partner over MCS to detect failure. DHCP Server does not utilize this method and it strictly relies on the failure notifications by MCS.

Note that the intercommunication link failure does not necessarily assume the same failed fate for the access links. In other words, it is perfectly possible (although unlikely) that both access links are operational while the inter-chassis communication link is broken.

The failure detection of the intercommunication link leads to certain failover state transitions on the DHCP server. The DHCP lease handling in the local-remote model depends on the particular failover state on the DHCP server and the duration of each failover state is determined by preconfigured timers.

3.3.1.5 DHCP Server Failover States

The DHCP server, when paired in redundant fashion, can transition through several states:

- TRANSITION
- SHUTDOWN
- INIT
- STARTUP — In this state, the DHCP server recovers leases from the MCS database and does not respond to any unicast and broadcast messages.
- NORMAL – In this state, the 7750 SR and 7450 ESS DHCP Server is serving all IP leases from the local and access-driven IP addresses-ranges and prefixes (assigning new leases and extending existing ones). Remote IP addresses-ranges and prefixes are not served (new or existing ones).
- PRENORMAL — In this state, the DHCP recovers leases from the MCS database after a resolved communication failure. The DHCP server responds to unicast and broadcast messages for addresses in the local ranges, and to unicast messages for addresses in the remote ranges.

- COMMUNICATION INTERUPTED – IP addresses and prefixes under the local and access-driven IP address-range and prefix are served in the same fashion as in the NORMAL state. The IP addresses and prefixes under the remote IP address-range and prefix are renewed. However, the new address and prefix from the remote address-range and prefix are not allocated until the partner-down timer expires, the failover state consequently transitions into PARTNER DOWN and the MCLT timer while in the PARTNER-DOWN state expires. This is necessary in case that the failure occurred only on the intercommunication link while the access link is still operational (DHCP server nodes become isolated).

  COMMUNICATION INTERUPTED state indicates that there is a failure of some kind, but it cannot be determined whether an entire node failed or only the inter-chassis link. The access layer may still be fully operational, but the new leases (including RENEWS/REBINDS) cannot be synchronized between the two peers.
- PARTNER DOWN – Once the DHCP Server reaches this state, the remote IP address-range and prefix is taken over and after an additional time period of MCLT (Maximum Client Lead Time), the new IP addresses and prefixes from it can be delegated to clients. PARTNER-DOWN state is an indication (and assumption at the same time) that the remote node is truly down.

  Otherwise, the IP address duplication may occur in case that all of the following conditions are met:
  - the DHCP server nodes become isolated
  - the failover state is PARTNER-DOWN
  - DHCP/PPPoX clients have simultaneous access to the same IP address-range and prefix on the both DHCP servers

3.3.1.6 Lease Time Synchronization

7750 SR and 7450 ESS DHCP server state synchronization is different from the lease state synchronization of the subscriber host itself.

- Each subscriber host lease state is synchronized via sub-mgmt client application. The host lease state can be seen via the output of following CLI command:

  show service id *id* **dhcp lease-state**

  The output of this command represents the state of our internal DHCP relay (client).
- The DHCP server lease states can be observed via the following CLI command:

  show routed *id* **dhcp local-dhcp-server** *name* **lease**

The concern is with the latter, the 7750 SR and 7450 ESS DHCP server lease state synchronization. To ensure un-interrupted IP lease renewal process after a failure in the network, the DHCP server lease time that is synchronized between the 7750 SR and 7450 ESS DHCP servers must always lead the currently assigned lease time for the period of the anticipated lease time in the next period.

For comparison purposes the two flow diagrams are juxtaposed in Figure 28:

- The right side does not include any lead time during synchronization
- The left side includes the lead time during synchronization

Figure 28 **Potential Expiration Time**

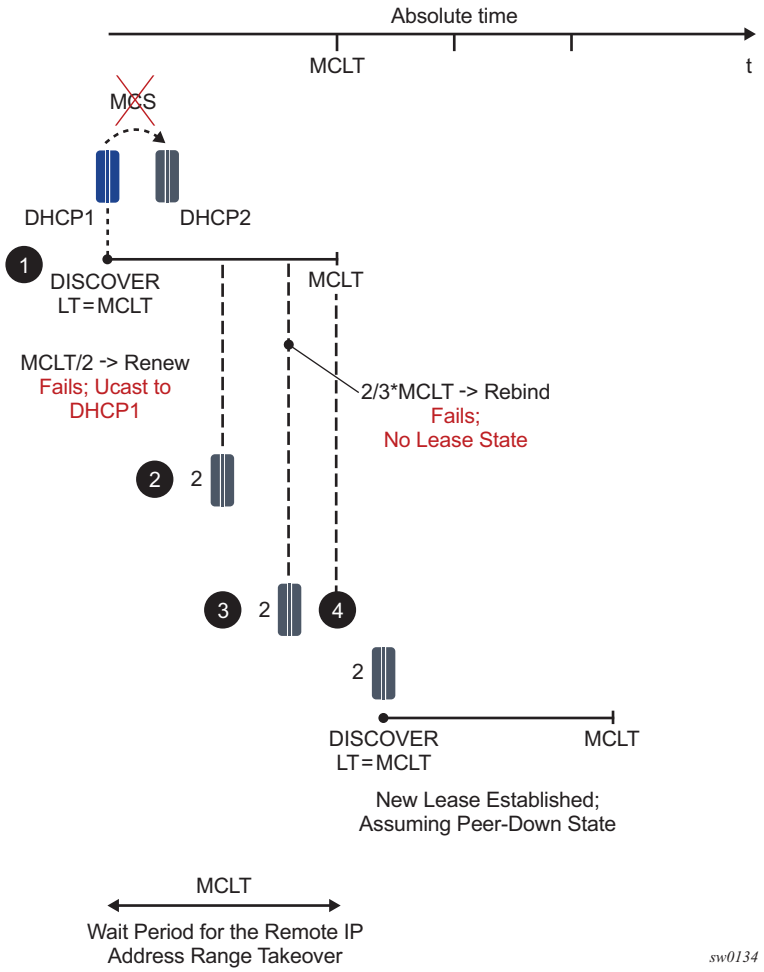

On the left side of the graph, the lease time is synchronized to the same value to which it was renewed.

In point 3, the primary DHCP server renews the lease time but fails to synchronize it due to its own failure (crash). The secondary server (2) has the old lease time A' in its database. The next time the client tries to REBIND its address and prefix, the lease in the secondary server expires (4). As a result, the IP address and prefix are not renewed.

To resolve this, the primary server must synchronize the lease time as the current RENEW time plus the next lease time. This way, as depicted in point 3, when the REBIND reaches server 2, the lease in its DB is active (4) and the server 2 can extend the lease for the client.

3.3.1.7 Maximum Client Lead Time (MCLT)

Maximum Client Lead Time (MCLT) is the maximum time that the 7750 SR and 7450 ESS DHCP Server can extend the lease time to its clients beyond the lease time currently know by the 7750 SR and 7450 ESS DHCP partner node. By default this time is relatively short (10 minutes).

The purpose of the MCLT is explained in the following scenario:

The local DHCP server assigns a new IP lease to the client but it crashes before it sends a sync update to the partner server. Because of the local DHCP server failure, the remote DHCP server is not aware of the IP address and prefix that was allocated on the local DHCP server. This condition creates the possibility that the remote DHCP server allocates the same address and prefix to another client. This would cause IP address and prefix duplication. MCLT is put in place to prevent this scenario.

MCLT based solution is shown in Figure 29.

Figure 29 **Address and Prefix Allocation without Synchronization**

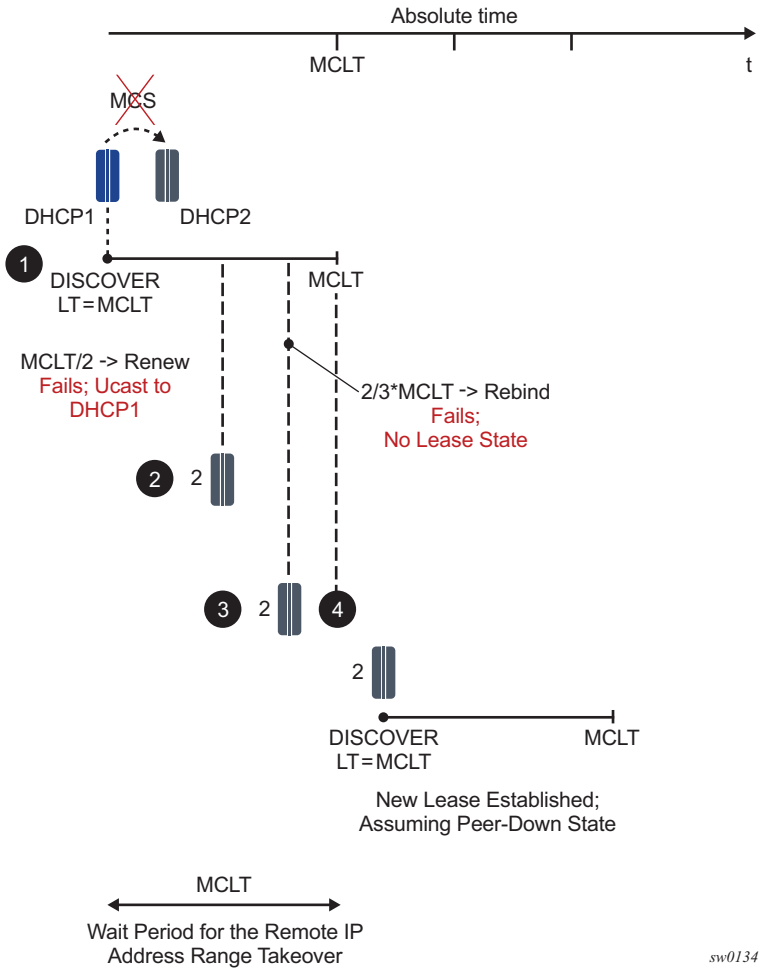

The sequence of events is the following:

1. DHCP Server 1 is the local DHCP server (with the local address-range and prefix) that creates the IP lease state for a new client. The initial lease-time assigned to the client is MCLT which is normally shorter than the requested lease time.

   This DHCP server fails before it gets a chance to synchronize the lease state with the DHCP Server 2 (remote DHCP server with the remote address-range and prefix).

2. The remote DHCP server transitions into the PARTNER-DOWN state (assuming that the partner-down timer is 0). In this state the remote DHCP server can extend the lease time to the existing clients but it can not assign a new lease for a period of MCLT. In MCLT/2 a new RENEW request is sent directly to the local DHCP server. This server is DOWN and therefore it cannot reply.
3. The client broadcasts a REBIND request that reaches the remote DHCP server. The remote DHCP server has no knowledge of the requested lease and therefore it does not reply.
4. The lease for the client expires and the client must reinitiate the IP address and prefix assignment process.

Since the remote DHCP server is not aware of the lease state that was assigned by the local DHCP server, there is a chance that the remote DHCP server assigns to the new client the same IP address and prefix already allocated by the local DHCP server just before it crashed. This is why the remote DHCP server needs to wait for the MCLT time to expire so that the IP addresses and prefixes allocated (but never synchronized) by the local DHCP server can time out.

When the communication channel between the chassis is interrupted, two scenarios are possible:

1. The entire node becomes unavailable. In this case the redundant node takes over and it starts reducing the lease time until the lease time reaches MCLT.
   - COMMUNICATION-INTERUPTED

     The remote DHCP server only renews the leases but does not delegate new ones (primary is DOWN).

     The local DHCP server renews the leases (which eventually trickle down to MCLT) and delegates the new ones with the lease time of MCLT (secondary is down).
   - PARTNER-DOWN

     The remote DHCP server starts delegating new IP addresses and prefixes from the remote address-range and prefix after MCLT (primary is down). The lease time of the new clients is MCLT. A lease cannot be assigned for a period longer than what is agreed with the peer incremented with the MCLT. As for a “new” lease nothing is agreed yet so the sum falls back to the MCLT itself.
2. The communication channel is down but the remote DHCP server is not (meaning that the clients have still access to both servers). The behavior in this case is following:
   - COMMUNNICATION-INTERUPTED

The remote DHCP server keeps renewing existing leases but it does not delegate the new leases. There is little chance this will happen as the clients continue to send RENEWS via unicast to the local DHCP server which is still active. The non-synched leases in the remote DHCP server times out. The local DHCP server starts ticking down the lease time to MCLT.

The local DHCP server continues to delegate the new leases, although with the MCLT lease time.

– PARTNER-DOWN State

The local DHCP server continues to extend the existing leases but also starts delegating new IP leases once the initial MCLT elapses.

Both local and remote DHCP server delegates new leases in the PARTNER-DOWN state (although the remote DHCP server in PARTNER-DOWN state must wait an additional MCLT period before it can start delegating new leases).

3.3.1.8 Sharing IPv4 Address-Range or IPv6 Prefixes

Access-driven DHCP server redundancy model ensures uninterrupted IP address assignment service from a single IP address-range and prefix if an access link forwards BNG fails. To avoid duplicate address allocation, there must be a single active path available from the clients to only one of the 7750 SR and 7450 ESS DHCP servers in redundant configuration. This single active path is ensured via a protection mechanism in dual-homed environment in the access side of the network. The supported protection mechanisms in the access in 7750 SR and 7450 ESS are SRRP or MC-LAG.

In access-driven DHCP redundancy model, the DHCP relay in each 7750 SR and 7450 ESS node must point only to the IP address of the local DHCP server. In other words, the DHCP messages received on one DHCP server should never be relayed to the other. Since the IP address-ranges and prefixes are shared between the DHCP servers, accessing both DHCP servers with the same DHCP request can cause DHCP lease duplication. Moreover, the IP addresses of both DHCP servers must be the same in both nodes. Otherwise the DHCP renew process would fail.

Granting new leases out of the shared IP address-ranges or prefixes that are configured as access-driven is not dependent on the state of the inter-chassis communication link (MCS). Instead, the new leases can be granted from both nodes simultaneously and it is the role of the protection mechanism in the access to ensure that a single path to either server is always active.

This model allows the newly active node, after a SRRP/MC-LAG switchover, to be able to serve new clients immediately from the same (shared) IP range or prefix. At the same time, upon the switchover, the corresponding subscriber-interface route is re-evaluated for advertisement with a higher routing metric to the network side via SRRP awareness. The end result is that by aligning the subscriber-interface routes or prefixes with access-driven DHCP address ranges or prefixes in DHCP server, the IP address ranges or prefixes are advertised to the network side only from the actively serving node (SRRP Master or active MC-LAG node). This is performed indirectly via the corresponding subscriber-interface route that is aligned (by configuration) with the DHCP address range or prefix in access-driven mode.

If SRRP or MC-LAG is not deployed in conjunction with the access-driven configuration option, the IP address duplication could occur.

Multi-chassis redundancy relies on MCS to synchronize various client applications. (subscriber states, DHCP states, IGMP states, and so on) between the two chassis. Therefore the links over which MCS peering session is established must be highly redundant. Failed MCS peering session renders dual-chassis redundancy non-operational. Considering this fact, the DHCP failover scenario with SRRP/MC-LAG and shared IP address range should be evaluated considering the following cases described in Table 39.

Table 39 **DHCP Failover Scenarios**

Access Related Failure	Inter-Chassis Link State	Number of Failures	Action
None	None	None	Only the MASTER SRRP BNG grants IP leases. The subnet tied to the SRRP instance is advertised to the network side on the MASTER SRRP node.
None	COMM-INT	Possibly multiple	Only the MASTER SRRP BNG grants IP leases. Communication link between the two chassis has failed and the DHCP states cannot be synchronized. Operator is required to restore the links between the chassis.
None	PARTNER-DOWN	Possibly multiple	Same as above. The premise of this deployment model in general (SRRP/MC-LAG + access-driven) is that there is only one path leading to the DHCP server active. This path is governed by SRRP. Therefore there is no change in behavior between the PARTNER-DOWN and COMM-INT states in this particular scenario.

Table 39 **DHCP Failover Scenarios (Continued)**

Access Related Failure	Inter-Chassis Link State	Number of Failures	Action
Access Link Towards the Master SRRP	NORMAL	Single	SRRP switches over. The new IP lease grants continues from new SRRP Master using the same IP address range. IP leases are synched OK. Subnets are advertised from new Master via SRRP awareness.
Access Link Towards the Master SRRP	COMM-INT	Multiple	SRRP switch over. The new leases can be handed from the DHCP server on the newly SRRP Master node. However, this DHCP server may not have its lease state table up to date since the inter-chassis communication link is non-operational. Consequently, new SRRP Master node may start handing out leases that are already allocated on the node with the failed link. In this case, IP address duplication would ensue. This is why it of utmost importance that the intercommunication chassis link is well protected and the only event that causes it to go down is when the entire node goes down. Otherwise the nodes becomes isolated from each other and synchronization becomes non-operational.
Access Link Towards the Master SRRP	PARTNER-DOWN	Multiple	Same as above.
Access Link Towards the Standby SRRP	NORMAL	Single	No effect on the operation since everything is active on the Master SRRP node anyway.
Access Link Towards the Standby SRRP	COMM-INT	Multiple	Intercommunication link is broken. The Master SRRP node continues handing out the new leases and renewing the old ones. However, they are not synchronized to the peering node.
Access Link Towards the Standby SRRP	PARTNER-DOWN	Multiple	Same as above.

Table 39 **DHCP Failover Scenarios (Continued)**

Access Related Failure	Inter-Chassis Link State	Number of Failures	Action
Entire Master Node	COMM-INT	Single	SRRP switches over. However, lease duplication may occur on the newly Master since the intercommunication link is broken and this newly SRRP Master is not aware of DHCP leases that the peer (failed node) may have allocated to the clients while the intercommunication-link was broken.
Entire Master Node	PARTNER-DOWN	Single	Same as above.
Entire Standby Node	COMM-INT	Single	Operation continues, but multi-chassis redundancy is lost.
Entire Standby Node	PARTNER-DOWN	Single	Same as above.

A possible deployment scenario is shown in Figure 30.

Figure 30 **Failover Scenario with SRRP and DHCP in Access-Driven Mode**

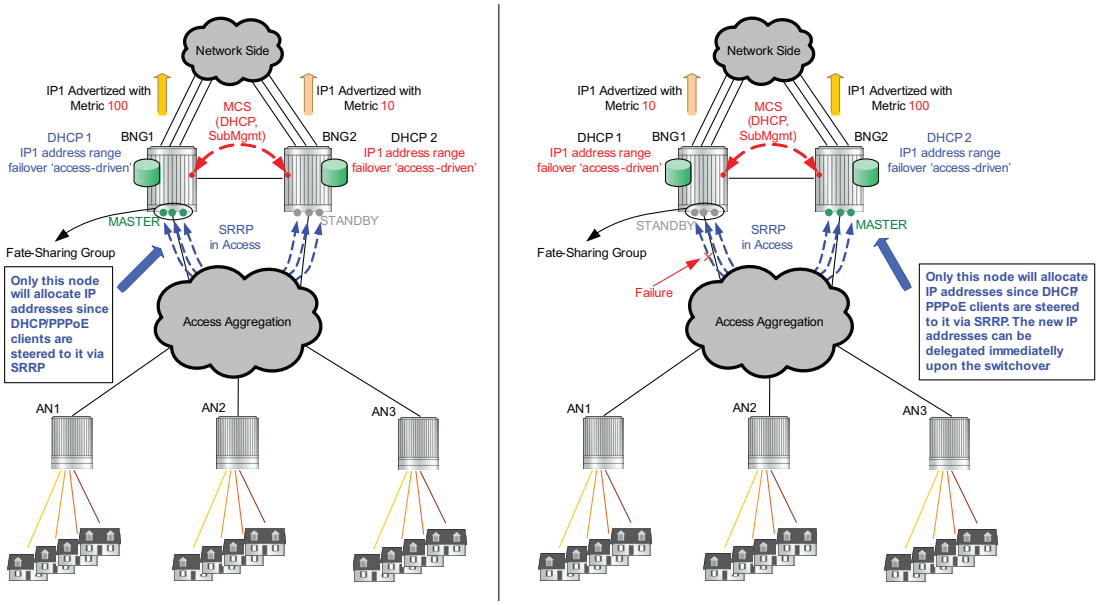

3.3.1.9 Fast-Switchover of IP Address and Prefix Delegation

Some deployments require that the remote IP address and prefix range starts delegating new IP addresses and prefixes upon the failure of the intercommunication link, without waiting for the intercommunication link to transition from the COMM-INT state into the PARTNER-DOWN state and the MCLT to expire while in PARTNER-DOWN state.

In other words, the takeover of the remote IP address-range and prefix should follow the failure of the intercommunication link, without any significant delays.

This can be achieved by configuring both of the following two items under the dhcp failover CLI hierarchy:

- The partner-down-delay must be set to 0. This causes the intercommunication link to bypass the COMM-INT state upon the failure and transition straight into the PARTNER-DOWN state. The remote IP address-range and prefix can be taken over only in PARTNER-DOWN state, once the MCLT expires.
- The ignore-mclt-on-takeover flag must be enabled. With this flag enabled, the remote IP address and prefix can be taken over immediately upon entering the PARTNER-DOWN state of the intercommunication link, without having to wait for the MCLT to expire. By setting this flag, the lease times of the existing DHCP clients, while the intercommunication link is in the PARTNER-DOWN state, are reduced to the MCLT over time and all new lease times are set to MCLT. This behavior remain the same as originally intended for MCLT. this functionality must be exercised with caution. One needs to keep in mind that the partner-down-delay and MCLT timers were originally introduced to prevent IP address duplication in cases where DHCP redundant nodes transition out-of-sync due to the failure of intercommunication link. These timers (partner-down-delay and MCLT) would ensure that during their duration, the new IP addresses and prefixes are delegated only from one node, the one with local IP address-range and prefix. The drawback is that the new IP address delegation is delayed and service is impacted.

If the intercommunication link could be guaranteed to always available, then the DHCP nodes would stay in sync and the two timers would not be needed. This is why it is importance that in this mode of operation, the intercommunication link is well protected by providing multiple paths between the two DHCP nodes. The only event that should cause intercommunication link to fail is the entire nodal failure. This failure is acceptable since in this case only one DHCP node is available to provide new IP addresses and prefixes.

3.3.1.9.1 DHCP Server Synchronization and Local PPPoX Pools

Since there is no classical DHCP lease state maintained for local PPPoX pools, the IP addresses are not synchronized via DHCP Server. Instead they are synchronized via PPPoX clients. Once the PPPoX subscriber is synchronized, the respective IP address lease is updated in the respective local pool.

For example:

- A PPPoE client is created on 7750 SR and 7450 ESS ‘A’.
- The IP address is assigned from the local poll on 7750 SR and 7450 ESS ‘A’.
- The PPPoE client is synchronized to the peering node 7750 SR and 7450 ESS ‘B’.
- Once the client is synchronized in the 7750 SR and 7450 ESS ‘B’, the IP address assignment is synchronized by the internal PPPoE process on the 7750 SR and 7450 ESS ‘B’ with the local pool.

One artifact of this behavior (IP address assignment in local DHCP pools is synchronized via PPPoX clients and not via DHCP server synchronization mechanism) is that during the node boot, the DHCP server must wait for the completion of PPPoX subscriber synchronization via MCS so that it learns which addresses and prefixes are already allocated on the peering node. Since the DHCP server can theoretically start assigning IP addresses before the PPPoX sync is completed, a duplicate address assignment my occur. For example an IP address lease can be granted via DHCP local pools while PPPoX sync is still in progress. Once the PPPoX sync is completed, the DHCP server may discover that the granted IP lease has already been allocated by the peering node. The most recent lease is kept and the other is removed from both systems. To prevent this scenario, a configurable timer can be set to an arbitrary value that renders sub-if non-operational until the timer expires. The purpose of this timer is to allow the PPPoX sync to complete before subscribers under the sub-intf can be served.

3.4 Local Address Assignment

3.4.1 Stateless Address Auto-Configuration

In the stateless auto-configuration model, hosts can be assigned address statically or dynamically. For static prefix assignment, LUDB and RADIUS can be used. For dynamic assignment, a pool name returned from LUDB or RADIUS and the local DHCPv6 server is used for address management. Although, the DHCPv6 server is used there are no lease time associated with the SLAAC prefix assigned to hosts. To utilize the local pool for SLAAC prefix assignment, the command **local-address-assignment** is used under group-interface. The client-application type **ppp-slaac** and/or **ipoe-slaac** must first be specified. Afterwards, the server name of the local-address-server must also be provisioned.

3.5 Configuring DHCP with CLI

This section provides information to configure DHCP using the command line interface.

3.5.1 Enabling DHCP Snooping

DHCP snooping is the process of copying DHCP packets and using the contained information for internal purposes. The BSA and BSR can use the snooped DHCP information to build anti-spoofing filters, populate the ARP table, send ARP replies, and so on.

For VPLS, DHCP snooping must be explicitly enabled (using the **snoop** command) on the SAP or SDP where DHCP messages ingress the VPLS instance. It is recommended to enable snooping on both the interface to the DHCP server (to snoop ACK messages) and the interface to the subscriber (to snoop RELEASE messages).

For IES and VPRN IP interfaces (VPRN is supported on the 7750 SR only), lease-populate enables DHCP snooping for the subnets defined under the IP interface. The number of allowed simultaneous DHCP sessions on a SAP or interface can be limited using the lease-populate command with the parameter number-of-entries specified. Enabling lease-populate and snoop commands is effectively enabling “standard subscriber management” as described in Standard and Enhanced Subscriber Management.

The following output displays an example of a partial BSA configuration with DHCP snooping enabled in a service:

```
*A:ALA-48>config>service# info
----------------------------------------------
...
        vpls 600 customer 701 create
            sap 1/1/4:100 split-horizon-group "DSL-group2" create
                description "SAP towards subscriber"
                dhcp
                    lease-populate 1
                    option
                        action replace
                        circuit-id
                        no remote-id
                    exit
                    no shutdown
                exit
            exit
            mesh-sdp 2:800 create
                dhcp
```

```
                    snoop
                exit
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-48>config>service#
```

3.5.2 Configuring Local User Database Parameters

A local user data base defines a collection of host entries. There are two types of hosts: PPP and IPoE. A local user database can be used to:

- Authenticate PPP clients. For this only the host entries configured in the **ppp** CLI are matched.
- Authenticate IPoE hosts (DHCPv4, DHCPv6 IA-NA/IA-PD, SLAAC). The host entries configured in the **ipoe** CLI context are matched.
- Perform authentication and address management for the local DHCPv4 server. For this, both PPP and IPoE sections can be used depending on the client type indicated by a vendor-specific sub-option inside Option 82 of the DHCPv4 message.

Each host can be identified by a set of values. However, at any point in time only four of these values are taken into account for IPoE as defined by the **ipoe match-list** option and only three are considered for PPP as defined in the **ppp match-list** option.

When trying to find a matching host entry, attempts are made to match as many items as possible. If several hosts match an incoming IPoE packet, the one with most match criteria is taken.

One host entry can map on several physical clients. For instance, when using a circuit ID, by masking when the interface ID is used, the host entry is used for all the clients on that same interface.

IPoE host identification includes:

- Circuit ID — this field also matches the DHCPv6 interface-id field
- MAC address
- Remote ID — matches on the remote-id sub-option in option 82 for DHCPv4 clients and on the remote-id option (including enterprise-id field) for DHCPv6 clients
- Option 60 from DHCPv4 message — only first 32 bytes are looked at

- SAP ID
- Service ID
- String from vendor-specific sub-option of Option 82
- System ID
- derived-id — a string provided via a DHCP Python script
- dual-stack-remote-id — matches on the remote-id sub-option in option 82 for DHCPv4 clients and on the remote-id field in the remote-id option (without enterprise-id) for DHCPv6 clients
- encap-tag-range — matches on VLAN tag ranges
- IP — matches on the source IPv4/IPv6 address of a data-trigger packet

PPP host identification includes:

- Circuit ID
- MAC address
- Remote id
- User name, either complete user name, domain part only, or host part only
- Derived ID — a string provided by Python script

When a host cannot be inserted in the lookup database, it is placed in an unmatched-hosts list. This can occur due to:

- Another host with the same host-identification exists. Only the host-identification that is specified in the match-list is taken into account for this.
- A host has no host-identification specified in the match-list.

When used for PPPoE-authentication, the fields are used as follows:

- password — Verifies the PPPoE user password. This is mandatory. If no password is required then it must be explicitly set to **ignore**.
- address:
  - no address — No address information. The address must be obtained by other means, either RADIUS or DHCP server.
  - gi-address — No meaning in this context. The address must be obtained by other means, either RADIUS or DHCP server.
  - use-pool-from-client — No meaning in this context, address must be obtained by other means, either RADIUS or DHCP server.
  - pool-name — The address must be obtained by other means, either RADIUS or a DHCP-server. When a DHCP server is used, this pool-name is included in Option 82 vendor-specific sub-option.
  - ip-address — This IP address is offered to the client.

- Identification-strings — Returns the strings used for enhanced subscriber management (ESM).
- Options — Only DNS servers and NBNS server are used, others are ignored.

When used from the DHCP server, the following applies:

- password — Not used.
- address — Defines how the address must be allocated for this host.
  - no address — The host is not allowed. The clients mapping to this host will not get an IP address.
  - gi-address — Finds the matching subnet and an IP address is taken from that subnet.
  - pool-name — A free IP address is taken from that pool.
  - ip-address — This address is offered to the client.
  - use-pool-from-client — Use the **poolname** in the Option 82 vendor-specific sub-option. If no **poolname** is provided there, falls back to the DHCP server default (none or use-gi-address).
- identification-strings — The operator can specify subscriber management strings and in which option the strings are sent back in dhcp-offer and dhcp-ack messages.
- options — The operator defines which options specific to this host should be sent back in the dhcp-offer and dhcp-ack messages. The options defined here override options defined on the pool-level and subnet-level inside the local DHCP server.

The circuit ID from PPPoE or from Option 82 in IPoE messages can be masked in following ways:

- prefix-length — Drop a fixed number of bytes at the beginning of the circuit-id.
- suffix-length— Drop a fixed number of bytes at the end of the circuit-id.
- prefix-string — The matching string is dropped from the beginning of the circuit-id. The matching string can contain wildcards (*). For example: incoming circuit-id mybox|3|my_interface|1/1/1:22 masked with *|*| will leave my_interface|1/1/1:22.
- suffix-string — The matching string is dropped at the end of the circuit-id. For example: incoming circuit-id mybox|3|my_interface|1/1/1:22 masked with |* will result in mybox|3|my_interface.

The following is an example of a local user database used for PPPoE authentication:

```
*A:ALA-48>config>subscr-mgmt# info
----------------------------------------------
...
        local-user-db "pppoe user db"
```

```
                description "pppoe authentication data base"
                ppp
                    match-list username circuit-id
                    mask prefix-string "*|*|" suffix-string "|*"
                    host "john" create
                        host-identification
                            username "john" no-domain
                        exit
                        password pap "23T8yPoe0w1R.BPGHB98i0qhJf7ZlZGCtXBKGnjrIrA" hash2
                        no shutdown
                    exit
                    host "test.com" create
                        host-identification
                            username "test.com" domain-only
                        exit
                        password ignore
                        no shutdown
                    exit
                    host "john@test.com" create
                        host-identification
                            username "john@test.com"
                        exit
                        password pap "23T8yPoe0w0Tlf1yCb4hskknvTYLqA2avvBB567g3eQ" hash2
                        identification-strings 122 create
                            subscriber-id "john@test.com"
                            sla-profile-string "sla prof1"
                            sub-profile-string "subscr profile 1"
                            ancp-string "ancp string"
                            inter-dest-id "inter dest"
                        exit
                        no shutdown
                    exit
                    host "john@test.com on interface group-if"
                        host-identification
                            circuit-id string "group-if"
                            username "john@test.com"
                        exit
                        password pap "23T8yPoe0w1R.BPGHB98i0qhJf7ZlZGCtXBKGnjrIrA" hash2
                        address 10.1.2.3
                        no shutdown
                     exit
                exit
                no shutdown
            exit
...
-------------------------------------------------
*A:ALA-48>config>subscr-mgmt#
```

The following are some examples when a user tries to set up PPPoE:

- john@test.com tries to setup PPPoE with circuit-id pe_23|3|group-if|1/1/1: host john@test.com on interface group-if will match, the PAP password is checked and the IP address 10.1.2.3 is given to PPPoE to use for this host.
- john@test.com (on another interface): host john@test.com will match, the PAP password is checked, and identification strings are returned to PPPoE.

- nokie@test.com: host test.com will match, no password check, the user is allowed.
- john@nokie.com: host john will match and the password is checked.
- anybody@anydomain: will not match and will not be allowed.

The following is an example of a local user database used for DHCP server for IPoE clients:

```
*A:ALA-50>config>subscr-mgmt# info
----------------------------------------------
...
        local-user-db "dhcp server user db"
            description "dhcp server user data base"
            ipoe
                match-list circuit-id mac
                mask prefix-string "*|*|" suffix-string "|*"
                host "mac 3 on interface" create
                    host-identification
                        circuit-id string "group-if"
                        mac 00:00:00:00:00:03
                    exit
                    address 10.0.0.1
                    no shutdown
                exit
                host "maskedCircId" create
                    host-identification
                        circuit-id string "group-if"
                    exit
                    address pool "pool 1"
                    identification-strings 122 create
                        subscriber-id "subscriber 1234"
                        sla-profile-string "sla prof 1"
                        sub-profile-string "sub prof 1"
                        ancp-string "ancpstring"
                        inter-dest-id "inter dest id 123"
                    exit
                    options
                        netbios-name-server 1.2.3.4
                        lease-time min 2
                    exit
                    no shutdown
                exit
            exit
            no shutdown
        exit
...
----------------------------------------------
*A:ALA-50>config>subscr-mgmt#
```

The following is an access example:

- MAC 00:00:00:00:00:03 on circuit-id pe5|3|group-if|1/1/1: host mac 3 on interface is matched and address 10.0.0.1 is offered to the IPoE client.

- Another MAC on circuit-id pe5|3|group-if|2/2/2: host maskedCircId is matched and an address is taken from pool1 (defined in the DHCP server). The identification-strings are copied to Option 122 in the dhcp-offer and dhcp-ack messages. The options defined here will also be copied into dhcp-offer and dhcp-ack messages.
- The circuit-id pe5|3|other_group_if|1/1/3: no host is matched. The client will only get an IP address if on DHCP server level you defined the use-gi-address parameter and the gi-address matches a subnet.

The following is an example of a local user database used for a DHCP server, only for PPPoE clients:

If PPPoE does not get an IP address from RADIUS or the local-user-db used for authentication, the internal dhcp-client is used to access a DHCP server which can be in the same node or in another node. These request are identified by inserting Option 82 sub-option client-id in the dhcp-discover and dhcp-request messages. When the DHCP server receives this request and has a user-db connected to it, then the PPPoE section of that user-db is accessed.

```
*A:ALA-60>config>subscr-mgmt# info
----------------------------------------------
...
        local-user-db "pppoe user db"
            description "pppoe authentication data base"
            ppp
                match-list username
                host "internet.be" create
                    host-identification
                        username "internet.com" domain-only
                    exit
                    address "pool_1"
                    no shutdown
                exit
                host "john@internet.com" create
                    host-identification
                        username "john@internet.com"
                    exit
                    identification-strings 122 create
                        subscriber-id "john@test.com"
                        sla-profile-string "sla prof1"
                        sub-profile-string "subscr profile 1"
                        ancp-string "ancp string"
                        inter-dest-id "inter dest"
                    exit
                    address use-gi
                    no shutdown
                exit
                host "malicious@internet.com"
                    host-identification
                        circuit-id string "group-if"
                        username "internet@test.com"
                    exit
                    no shutdown
                 exit
```

```
            exit
            no shutdown
        exit
...
-------------------------------------------
*A:ALA-60>config>subscr-mgmt#
```

The following is an access example:

- john@internet.com: GI is used to find a subnet and a free address is allocated form that subnet. Identification strings are returned in Option 122.
- anybody@internet.com: pool_1 is used to find a free IP address.
- malicious@internet.com: no address is defined. This user will not get an IP address.

The following is an example of associating a local user database to PPPoE for authentication for the 7750 SR.

```
A:pe5>config>service>vprn#
-------------------------------------------
            subscriber-interface "tomylinux" create
               address 10.2.2.2/16
               group-interface "grp_pppoe3" create
                   pppoe
                       e "pppoe"
                   exit
               exit
-------------------------------------------
A:pe5>config>service>vprn#
```

The following is an example of associating a local user database to a local DHCP server.

```
A:pe7>config>router>dhcp#
-------------------------------------------
                local-dhcp-server my_server
                    description "my dhcp server"
                    user-db "data base 1"
                        ...
                exit
-------------------------------------------
A:pe7>config>router>dhcp#
```

In PPPoE access scenarios without access node or with access nodes that do not insert PPPoE vendor specific tags Circuit-ID and/or Remote-ID, it may be required to configure this information in the local user database so that they can be picked up in pre-authentication phase and used for RADIUS authentication and reporting in RADIUS accounting messages. For example:

```
config>subscr-mgmt
```

```
local-user-db "ludb-1" create
    ppp
        match-list username
        host "host-1" create
            access-loop-information
                circuit-id string "LUDB inserted circuit-id"
                remote-id string "LUDB inserted remote-id"
            exit
            host-identification
                username "cpe-1@domain1.com"
            exit
            auth-policy "auth-policy-1"
            password ignore
            no shutdown
        exit
    exit
```

With PPPoE, when the system accesses a LUDB during a discovery phase, a matched host could return a second LUDB via a **user-db** configuration under the LUDB host context. This second database is accessed again during the PAP/CHAP phase. The following is an example:

```
local-user-db "padi-db" create
          ppp
              match-list derived-id
              host "testuser" create
                  host-identification
                      derived-id "testuser"
                  exit
                  msap-defaults
                      group-interface "g1"
                      service 500
                  exit
                  user-db "chap-db"
                  no shutdown
              exit
          exit
          no shutdown
      exit
      local-user-db "chap-db" create
          ppp
              match-list derived-id username
              host "testuser" create
                  host-identification
                      derived-id "testuser"
                      username "testuser"
                  exit
                  password chap "cYhRmQYWOkLW3s0LrtEnBjWlAwFa/1Kx" hash2
                  identification-strings 254 create
                      sla-profile-string "sla-2"
                  exit
                  no shutdown
              exit
          exit
          no shutdown
      exit
```

3.5.3 Configuring Option 82 Handling

Option 82, or the Relay Information Option is a field in DHCP messages used to identify the subscriber. The Option 82 field can already be filled in when a DHCP message is received at the router, or it can be empty. If the field is empty, the router should add identifying information (circuit ID, remote ID or both). If the field is not empty, the router can decide to replace it.

The following output displays an example of a partial BSA configuration with Option 82 adding on a VPLS service. Snooping must be enabled explicitly on a SAP.

```
A:ALA-1>config>service>vpls#
------------------------------------------------
            no shutdown
            description "Default tls description for service id 1"
            sap 1/1/11 split-horizon-group "2dslam" create
                dhcp
                    no description
                    snoop
                    no lease-populate
                    option
                        action replace
                        circuit-id ascii-tuple
                        no remote-id
                    exit
                    no shutdown
                exit
            exit
------------------------------------------------
A:ALA-1>config>service>vpls#
```

3.5.4 Enabling DHCP Relay

Lease populate and DHCP relay are different features in which are not both required to be enabled at the same time. DHCP relay can be performed without populating lease tables.

The following example displays DHCP relay configured on an IES interface:

```
A:ALA-48>config>service>ies>if# info
------------------------------------------------
                address 10.10.42.41/24
                local-proxy-arp
                proxy-arp
                    policy-statement "ProxyARP"
                exit
                sap 1/1/7:0 create
                    anti-spoof ip
                exit
                arp-populate
```

```
                    dhcp
                        description "relay_ISP1"
                        server 10.200.10.10 10.200.10.20
                        lease-populate 1
                        no shutdown
                    exit
----------------------------------------------
A:ALA-48>config>service>ies>if#
```

3.6 Triple Play DHCP Command Reference

3.6.1 Command Hierarchies

Note:

- The command trees in this section are limited to those commands specific to Triple Play services. For the full command trees for a specific service type, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* and the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN*.

- System DHCP Commands
- DHCP Filter Commands
- Router DHCP Local User Database Commands
- Local User Database Commands
- Service DHCP Commands
- Show Commands
- Tools Commands
- Clear Commands
- Debug Commands

3.6.1.1 System DHCP Commands

```
config
    — system
        — dhcp6
            — adv-noaddrs-global [esm-proxy] [esm-relay] [relay] [server]
            — no adv-noaddrs-global
```

3.6.1.2 DHCP Filter Commands

```
config
    — filter
        — [no] dhcp-filter filter-id [create]
            — default-action bypass-host-creation
            — default-action drop
```

```
            — no default-action
            — description description-string
            — no description
            — entry entry-id [create]
            — no entry entry-id
                — action bypass-host-creation
                — action drop
                — no action
                — option dhcp-option-number {present | absent}
                — option dhcp-option-number match hex hex-string [exact] [invert-
                    match]
                — option dhcp-option-number match string ascii-string [exact] [invert-
                    match]

config
    — filter
        — [no] dhcp6-filter filter-id [create]
            — default-action bypass-host-creation [na] [pd]
            — default-action drop
            — no default-action
            — description description-string
            — no description
            — entry entry-id [create]
            — no entry entry-id
                — action bypass-host-creation [na] [pd]
                — action drop
                — no action
                — option dhcp6-option-number {present | absent}
                — option dhcp6-option-number match hex hex-string [exact] [invert-
                    match]
                — option dhcp6-option-number match string ascii-string [exact] [invert-
                    match]
```

3.6.1.3 Router DHCP Local User Database Commands

```
config
    — router
        — dhcp
        — dhcp6
            — local-dhcp-server server-name [create]
            — no local-dhcp-server server-name
                — [no] allow-lease-query
                — description description-string
                — no description
                — failover
                    — [no] ignore-mclt-on-takeover
                    — maximum-client-lead-time [maximum-client-lead-time]]
                    — no maximum-client-lead-time
                    — partner-down-delay [partner-down-delay]
                    — no partner-down-delay
```

- **peer** *ip-address* **tag** *sync-tag*
- **no peer**
- [**no**] **shutdown**
- [**no**] **startup-wait-time** [**min** *minutes*] [**sec** *seconds*] [**hrs** *hours*]

- [**no**] **force-renews**
- [**no**] **ignore-rapid-commit**
- [**no**] **interface-id-mapping**
- **lease-hold-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
- **no lease-hold-time**
- [**no**] **lease-hold-time-for**
  - [**no**] **internal-lease-ipsec**
  - [**no**] **solicited-release**
- **pool** *pool-name* [**create**]
- **no pool** *pool-name*
  - **delegated-prefix-length** *bits* [**minimum** *prefix-length*] [**maximum** *prefix-length*]
  - **no delegated-prefix-length**
  - **description** *description-string*
  - **no description**
  - [**no**] **exclude-prefix** *ipv6-prefix/prefix-length*
  - **failover**
    - [**no**] **ignore-mclt-on-takeover**
    - **maximum-client-lead-time** [*maximum-client-lead-time*]
    - **no maximum-client-lead-time**
    - **partner-down-delay** [*partner-down-delay*]
    - **no partner-down-delay**
    - **peer** *ip-address* **tag** *sync-tag*
    - **no peer**
    - [**no**] **shutdown**
    - [**no**] **startup-wait-time** [**min** *minutes*] [**sec** *seconds*] [**hrs** *hours*]
  - **max-lease-time** [*max-lease-time*]
  - **no max-lease-time**
  - **min-lease-time** [*min-lease-time*]
  - **no min-lease-time**
  - **minimum-free** *minimum-free* [**percent**] [**event-when-depleted**]
  - **no minimum-free**
  - [**no**] **nak-non-matching-subnet**
  - **offer-time** [**min** *minutes*] [**sec** *seconds*]
  - **no offer-time**
  - **options**
    - **custom-option** *option-number* **address** [*ip-address*]
    - **custom-option** *option-number* **domain** *domain-string*
    - **custom-option** *option-number* **hex** *hex-string*
    - **custom-option** *option-number* **string** *ascii-string*
    - **no custom-option** *option-number*
    - **dns-server** *ip-address* [*ip-address*]
    - **no dns-server**
    - **domain-name** *domain-name*
    - **no domain-name**
    - **lease-rebind-time** [*lease-rebind-time*]
    - **no lease-rebind-time**
    - **lease-renew-time** [*lease-renew-time*]
    - **no lease-renew-time**

- **lease-time** [*lease-time*]
- **no lease-time**
- **netbios-name-server** *ip-address* [*ip-address*]
- **no netbios-name-server**
- **netbios-node-type** *netbios-node-type*
- **no netbios-node-type**

— **prefix** *ipv6-addr/prefix-len* [**failover** {**local** | **remote** | **access-driven**}] [**pd**] [**wan-host**] [**create**]

— **no prefix** *ipv6-addr/prefix-len*

- [**no**] **drain**
- **options**
  - **custom-option** *option-number* **address** *ipv6-address* [*ipv6-address*]
  - **custom-option** *option-number* **domain** *domain-string*
  - **custom-option** *option-number* **hex** *hex-string*
  - **custom-option** *option-number* **string** *ascii-string*
  - **no custom-option** *option-number*
  - **dns-server** *ipv6-address* [*ipv6-address*]
  - **no dns-server**
  - **domain-name** *domain-name*
  - **no domain-name**
- **preferred-lifetime** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
- **no preferred-lifetime**
- **rebind-timer** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
- **no rebind-timer**
- **renew-timer** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
- **no renew-timer**
- **thresholds**
  - [**no**] **minimum-free** **prefix-length** [*prefix-length*]
    - [**no**] **depleted-event**
    - **minimum** **percent** [*percent*] [**number** [*number*]]
    - **no minimum**
  - **valid-lifetime** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
  - **no valid-lifetime**

— **subnet** {*ip-address*/*mask* | *ip-address netmask*} [**create**]

— **no subnet** {*ip-address*/*mask* | *ip-address netmask*}

- [**no**] **address-range** *start-ip-address end-ip-address* [**failover** {**local** | **remote**| **access-driven**}]
- [**no**] **drain**
- [**no**] **exclude-addresses** *start-ip-address* [*end-ip-address*]
- **maximum-declined** *maximum-declined*
- **no maximum-declined**
- **minimum-free** *minimum-free* [**percent**] [**event-when-depleted**]
- **no minimum-free**
- **options**

```
                    — custom-option option-number address [ip-address]
                    — custom-option option-number hex hex-string
                    — custom-option option-number string ascii-string
                    — no custom-option option-number
                    — default-router ip-address [ip-address]
                    — no default-router
                    — subnet-mask ip-address
                    — no subnet-mask
                — thresholds
                    — [no] minimum-free prefix-length [prefix-length]
                        — [no] depleted-event
                        — minimum percent [percent]
                        — no minimum
            — use-gi-address [scope scope]
            — no use-gi-address
            — use-pool-from-client delimiter delimiter
            — use-pool-from-client
            — no use-pool-from-client
            — use-link-address [scope scope]
            — no use-link-address
            — user-db local-user-db-name
            — no user-db
            — user-ident user-ident
            — no user-ident
            — server-id duid-en hex hex-string
            — server-id duid-en string ascii-string
            — server-id duid-ll
            — no server-id
```

3.6.1.4 Local User Database Commands

- IPoE Commands
- PPP Commands

3.6.1.4.1 IPoE Commands

```
config
    — subscriber-mgmt
        — local-user-db local-user-db-name [create]
        — no local-user-db local-user-db-name
            — description description-string
            — no description
            — ipoe
                — host host-name [create]
                — no host host-name
                    — acct-policy acct-policy-name [duplicate acct-policy-name]
                    — no acct-policy
```

- **address gi-address** [**scope** *scope*]
- **address** *ip-address*
- **address pool** *pool-name* [**secondary-pool** *sec-pool-name*] [**delimiter** *delimiter*]
- **address use-pool-from-client** [**delimiter** *delimiter*]
- **no address**
- **auth-domain-name** *domain-name*
- **no auth-domain-name**
- **auth-policy** *policy-name*
- **no auth-policy**
- **diameter-application-policy** *policy-name*
- **no diameter-application-policy**
- **diameter-auth-policy** *name*
- **no diameter-auth-policy**
- **gi-address** *ip-address*
- **no gi-address**
- **host-identification**
  - **circuit-id string** *ascii-string*
  - **circuit-id hex** *hex-string*
  - **no circuit-id**
  - **derived-id** *derived-id-string*
  - **no derived-id**
  - **encap-tag-range start-tag** *start-tag* **end-tag** *end-tag*
  - **no encap-tag-range**
  - **ip-prefix** *ip-prefixlip-prefix-length*
  - **no ip-prefix**
  - **mac** *ieee-address*
  - **no mac**
  - **option60 hex** *hex-string*
  - **option60 string** *ascii-string*
  - **no option60**
  - **remote-id hex** *hex-string*
  - **remote-id string** *ascii-string*
  - **no remote-id**
  - **sap-id** *sap-id*
  - **no sap-id**
  - **service-id** *service-id*
  - **no service-id**
  - **string** *string*
  - **no string**
  - **system-id** *system-id*
  - **no system-id**
- **identification-strings** *option-number* [**create**]
- **no identification-strings**
  - **ancp-string** *ancp-string*
  - **no ancp-string**
  - **app-profile-string** *app-profile-string*
  - **no app-profile-string**
  - **category-map-name** *category-map-name*
  - **no category-map-name**
  - **inter-dest-id** *intermediate-destination-id*
  - **no inter-dest-id**
  - **sla-profile-string** *sla-profile-string*
  - **no sla-profile-string**

- — **spi-sharing-group-id** *group-id*
- — **no spi-sharing-group-id**
- — **sub-profile-string***sub-profile-string*
- — **no sub-profile-string**
- — **subscriber-id** *sub-ident-string*
- — **no subscriber-id**

— **interface** *ip-int-name* **service-id** *service-id*
— **ipv6-address** *ipv6-address*
— **no ipv6-address**
— **ipv6-delegated-prefix** *ipv6-prefix/prefix-length*
— **no ipv6-delegated-prefix**
— **ipv6-delegated-prefix-length** *bits*
— **no ipv6-delegated-prefix-length**
— **ipv6-delegated-prefix-pool** *pool-name*
— **no ipv6-delegated-prefix-pool**
— [**no**] **ipv6-lease-times**

- — **preferred-lifetime** [*preferred-lifetime*]
- — **preferred-lifetime infinite**
- — **no preferred-lifetime**
- — **rebind-timer** [*rebind-timer*]
- — **no rebind-timer**
- — **renew-timer** [*renew-timer*]
- — **no renew-timer**
- — **valid-lifetime** [*valid-lifetime*]
- — **valid-lifetime infinite**
- — **no valid-lifetime**

— **ipv6-slaac-prefix** *ipv6-prefix/prefix-length*
— **no ipv6-slaac-prefix**
— **ipv6-slaac-prefix-pool** *pool*
— **no ipv6-slaac-prefix-pool**
— **ipv6-wan-address-pool** *pool-name*
— **no ipv6-wan-address-pool**
— **link-address** *ipv6-address*
— **no link-address**
— **match-radius-proxy-cache**

- — **fail-action** {**continue** | **drop**}
- — **no fail-action**
- — **mac-format** *mac-format*
- — **no mac-format**
- — **match** {**circuit-id** | **mac** | **remote-id**}
- — **match option** [*number*] [**option6** [*number*]]
- — **match option6** [*number*]
- — **no match**
- — **server** [**service** *service-id*] **name** *server-name*
- — **no server**

— **mld-parameters**

- — **import** *policy-name*
- — **no import**

— **msap-defaults**

- — **group-interface** *ip-int-name* [**prefix** {*port-id*}]
- — **group-interface** *ip-int-name* [**suffix** {*port-id*}]
- — **no group-interface**
- — **policy** *msap-policy-name*
- — **no policy**

```
        — service service-id
        — no service
    — options
        — custom-option option-number address [ip-address]
        — custom-option option-number hex hex-string
        — custom-option option-number string ascii-string
        — no custom-option option-number
        — default-router ip-address [ip-address]
        — no default-router
        — dns-server address [ip-address]
        — no dns-server
        — domain-name domain-name
        — no domain-name
        — lease-rebind-time [lease-rebind-time]
        — no lease-rebind-time
        — lease-renew-time [lease-renew-time]
        — no lease-renew-time
        — lease-time [lease-time]
        — no lease-time
        — netbios-name-server ip-address [ip-address]
        — no netbios-name-server
        — netbios-node-type netbios-node-type
        — no netbios-node-type
        — subnet-mask ip-address
        — no subnet-mask
    — options6
        — dns-server ipv6-address [ipv6-address]
        — no dns-server
    — retail-service-id service-id
    — no retail-service-id
    — router-advertisement-policy policy
    — no router-advertisement-policy
    — server ip-address
    — no server
    — server6 ipv6-address
    — no server6
    — [no] shutdown
    — to-client-options
        — ipv4
            — option option-number address [ip-address]
            — option option-number hex hex-string
            — option option-number string ascii-string
            — no option option-number
        — ipv6
            — option option-number address [ip-address.]
            — option option-number hex hex-string
            — option option-number string ascii-string
            — option option-number domain domain-name
            — no option option-number
    — [no] to-server-options
        — ipv6
            — option option number address [ip-address]
            — option option-number hex hex-string
            — option option-number string ascii-string
```

```
                                    — option option-number domain domain-name
                                    — no option option-number
                            — wpp
                                    — initial-app-profile app-profile-name
                                    — no initial-app-profile
                                    — initial-sla-profile sla-profile-name
                                    — no initial-sla-profile
                                    — initial-sub-profile sub-profile-name
                                    — no initial-sub-profile
                                    — [no] portal router router-instance name wpp-portal-
                                        name
                                    — portal-group portal-group-name
                                    — no portal-group
                                    — [no] restore-disconnected
                    — mask type ipoe-match-type {[prefix-string prefix-string | prefix-length
                        prefix-length] [suffix-string suffix-string | suffix-length suffix-
                        length]}
                    — no mask type ipoe-match-type
                    — match-list match-type-1 [match-type-2]
                    — no match-list
```

3.6.1.4.2 PPP Commands

```
config
    — subscriber-mgmt
            — local-user-db local-user-db-name [create]
            — no local-user-db local-user-db-name
                    — description description-string
                    — no description
                    — ppp
                            — host host-name [create]
                            — no host host-name
                                    — [no] access-loop-encapsulation
                                            — encap-offset [type type]
                                            — no encap-offset
                                            — rate-down rate
                                            — no rate-down
                                    — access-loop-information
                                            — circuit-id sap-id
                                            — circuit-id string ASCII string
                                            — no circuit-id
                                            — remote-id mac
                                            — remote-id string ASCII string
                                            — no remote-id
                                    — acct-policy acct-policy-name [duplicate acct-policy-name]
                                    — no acct-policy
                                    — address gi-address [scope scope]
                                    — address ip-address[/prefix-length]
                                    — address pool pool-name [secondary-pool sec-pool-name]
                                        [delimiter delimiter]
                                    — address use-pool-from-client [delimiter delimiter]
```

- **no address**
- **auth-policy** *policy-name*
- **no auth-policy**
- **diameter-application-policy** *policy-name*
- **no diameter-application-policy**
- **diameter-auth-policy** *name*
- **no diameter-auth-policy**
- [**no**] **force-ipv6cp**
- **host-identification**
  - **circuit-id string** *ascii-string*
  - **circuit-id hex** *hex-string*
  - **no circuit-id**
  - **derived-id** *derived id-string*
  - **no derived-id**
  - **encap-tag-range start-tag** *start-tag* **end-tag** *end-tag*
  - **no encap-tag-range**
  - **mac** *ieee-address*
  - **no mac**
  - **remote-id hex** *hex-string*
  - **remote-id string** *ascii-string*
  - **no remote-id**
  - **sap-id** *sap-id*
  - **no sap-id**
  - **service-name** *service-name*
  - **no service-name**
  - **username** *user-name*
  - **username** *user-name* **no-domain**
  - **username** *user-name* **domain-only**
  - **no username**
- **identification-strings** *option-number* [**create**]
- **no identification-strings**
  - **ancp-string** *ancp-string*
  - **no ancp-string**
  - **app-profile-string** *app-profile-string*
  - **no app-profile-string**
  - **category-map-name** *category-map-name*
  - **no category-map-name**
  - **inter-dest-id** *intermediate-destination-id*
  - **no inter-dest-id**
  - **sla-profile-string** *sla-profile-string*
  - **no sla-profile-string**
  - **spi-sharing-group-id** *group-id*
  - **no spi-sharing-group-id**
  - **sub-profile-string** *sub-profile-string*
  - **no sub-profile-string**
  - **subscriber-id** *sub-ident-string*
  - **no subscriber-id**
- [**no**] **ignore-df-bit**
- **interface** *ip-int-name* **service-id** *service-id*
- **no interface**
- **ipv6-address** *ipv6-address*
- **no ipv6-address**
- **ipv6-delegated-prefix** *ipv6-prefix/prefix-length*
- **no ipv6-delegated-prefix**

- **ipv6-delegated-prefix-length** *bits*
- **no ipv6-delegated-prefix-length**
- **ipv6-delegated-prefix-pool** *pool-name*
- **no ipv6-delegated-prefix-pool**
- [**no**] **ipv6-lease-times**
  - **preferred-lifetime** [*preferred-lifetime*]
  - **preferred-lifetime infinite**
  - **no preferred-lifetime**
  - **rebind-timer** [*rebind-time*]
  - **no rebind-timer**
  - **renew-timer** [*renew-timer*]
  - **no renew-timer**
  - **valid-lifetime** [*valid-lifetime*]
  - **valid-lifetime infinite**
  - **no valid-lifetime**
- **ipv6-slaac-prefix** *ipv6-prefix/prefix-length*
- **no ipv6-slaac-prefix**
- **ipv6-slaac-prefix-pool** *pool*
- **no ipv6-slaac-prefix-pool**
- **ipv6-wan-address-pool** *pool-name*
- **no ipv6-wan-address-pool**
- **l2tp**
  - **group** *tunnel-group-name* [**service-id** *service-id*]
  - **no group**
- **mld-parameters**
  - **import** *policy-name*
  - **no import**
- **msap-defaults**
  - **group-interface** *ip-int-name* [**prefix** {*port-id*}]
  - **group-interface** *ip-int-name* [**suffix** {*port-id*}]
  - **no group-interface**
  - **policy** *msap-policy-name*
  - **no policy**
  - **service** *service-id*
  - **no service**
- **options**
  - **custom-option** *option-number* **address** [*ip-address*]
  - **custom-option** *option-number* **hex** *hex-string*
  - **custom-option** *option-number* **string** *ascii-string*
  - **no custom-option** *option-number*
  - **dns-server** *address* [*ip-address*]
  - **no dns-server**
  - **netbios-name-server** *ip-address* [*ip-address*]
  - **no netbios-name-server**
- **options6**
  - **dns-server** *ipv6-address* [*ipv6-address*]
  - **no dns-server**
- **padi-auth-policy** *policy-name*
- **no padi-auth-policy**
- **pado-delay** *deci-seconds*
- **no pado-delay**
- **password ignore**
- **password** {**chap** *password* |**pap** *password*} [**hash**|**hash2**]
- **no password**

```
— pre-auth-policy policy-name
— no pre-auth-policy
— retail-service-id service-id
— router-advertisement-policy policy
— no router-advertisement-policy
— no retail-service-id
— [no] shutdown
— steering-profile steering-profile-name
— no steering-profile
— to-client-options
        — ipv6
                — option option-number address [ip-address]
                — option option-number hex hex-string
                — option option-number string ascii-string
                — option option-number domain domain-name
                — no option option-number
— mask type ppp-match-type {[prefix-string prefix-string | prefix-
        length prefix-length] [suffix-string suffix-string | suffix-
        length suffix-length]}
— no mask type ppp-match-type
— match-list match-type-1 [match-type-2]
— no match-list
```

3.6.1.5 Service DHCP Commands

3.6.1.5.1 IES DHCP Commands

```
config
    — service
        — ies service-id [customer customer-id] [create] [vpn vpn-id] [name name]
        — no ies service-id
            — interface ip-int-name [create]
            — interface ip-int-name [create] tunnel
            — no interface ip-int-name
                — dhcp
                    — description description-string
                    — no description
                    — gi-address ip-address [src-ip-addr]
                    — no gi-address
                    — lease-populate nbr-of-leases
                    — no lease-populate
                    — [no] option
                        — action {replace | drop | keep}
                        — no action
                        — circuit-id [ascii-tuple | ifindex | sap-id |vlan-ascii-
                            tuple]
                        — no circuit-id
                        — remote-id [mac | string string]
                        — no remote-id
```

```
                    — [no] vendor-specific-option
                        — [no] client-mac-address
                        — [no] pool-name
                        — [no] sap-id
                        — [no] service-id
                        — string text
                        — no string
                        — [no] system-id
                — proxy-server
                    — emulated-server ip-address
                    — no emulated-server
                    — lease-time [days days] [hrs hours] [min minutes] [sec
                        seconds] [override]
                    — no lease-time
                    — [no] shutdown
                — [no] relay-plain-bootp
                — relay-proxy [release-update-src-ip] [siaddr-override ip-
                    address]
                — no relay-proxy
                — python-policy name
                — no python-policy
                — server server1 [server2]
                — no server
                — [no] shutdown
                — [no] trusted
            — [no] ipv6
                — address ipv6-address/prefix-length [eui-64]
                — no address ipv6-address/prefix-length
                — [no] dhcp6-relay
                    — description description-string
                    — no description
                    — dhcp6-relay [nbr-of-leases]
                    — lease-populate [nbr-of-leases] route-populate [pd] na
                        [ta]
                    — lease-populate [nbr-of-leases] route-populate pd [na]
                        [ta] [exclude]
                    — lease-populate [nbr-of-leases] route-populate [pd] [na]
                    — no lease-populate
                    — link-address ipv6-address
                    — no link-address
                    — [no] neighbor-resolution
                    — [no] option
                        — interface-id
                        — interface-id [ascii-tuple]
                        — interface-id ifindex
                        — interface-id sap-id
                        — interface-id string
                        — no interface-id
                        — [no] remote-id
                    — python-policy name
                    — no python-policy
                    — server ipv6z-address [ipv6z-address]
                    — no server [ipv6z-address]
                    — source-address ipv6-address
```

```
                    — no source-address
                    — user-db local-user-db-name
                    — no user-db
                    — [no] use-arp
                    — [no] shutdown
                — [no] dhcp6-server
                    — max-nbr-of-leases max-nbr-of-leases
                    — no max-nbr-of-leases
                    — [no] prefix-delegation
                        — [no] prefix ipv6-address/prefix-length
                            — duid duid [iaid iaid]
                            — no duid
                            — preferred-lifetime seconds
                            — preferred-lifetime infinite
                            — no preferred-lifetime
                            — valid-lifetime seconds
                            — valid-lifetime infinite
                            — no valid-lifetime
                        — [no] shutdown
        — [no] subscriber-interface ip-int-name [create]
            — dhcp
                — client-applications {[dhcp] [ppp]}
                — no client-applications
                — description description-string
                — no description
                — gi-address ip-address [src-ip-addr]
                — no gi-address
                — python-policy name
                — no python-policy
                — relay-proxy [release-update-src-ip] [siaddr-override ip-
                    address]
                — no relay-proxy
                — server server1 [server2]
                — no server
                — [no] shutdown
                — [no] virtual-subnet
            — [no] group-interface ip-int-name [create]
                — dhcp
                    — client-applications {[dhcp] [ppp]}
                    — no client-applications
                    — description description-string
                    — no description
                    — [no] filter filter-id
                    — gi-address ip-address [src-ip-addr]
                    — no gi-address ip-address
                    — lease-populate [nbr-of-leases]
                    — lease-populate [nbr-of-leases] l2-header [mac ieee-
                        address]
                    — no lease-populate
                    — [no] match-circuit-id
                    — [no] option
                        — action {replace | drop | keep}
                        — no action
```

```
                    — circuit-id [ascii-tuple | ifindex | sap-id | vlan-
                        ascii-tuple]
                    — no circuit-id
                    — remote-id [mac | string string]
                    — no remote-id
                    — [no] vendor-specific-option
                        — [no] client-mac-address
                        — [no] pool-name
                        — [no] sap-id
                        — [no] service-id
                        — string text
                        — no string
                        — [no] system-id
                — proxy-server
                    — emulated-server ip-address
                    — no emulated-server
                    — lease-time [days days] [hrs hours] [min minutes]
                        [sec seconds] [override]
                    — no lease-time
                    — [no] shutdown
                — python-policy name
                — no python-policy
                — relay-proxy [release-update-src-ip] [siaddr-override
                    ip-address]
                — no relay-proxy
                — server server1 [server2]
                — no server
                — [no] shutdown
                — [no] trusted
                — user-db local-user-db-name
                — no user-db
```

3.6.1.5.2 VPLS DHCP Commands

```
config
    — service
        — vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [b-vpls | i-vpls]
            [etree] [name name]
        — no vpls service-id
            — sap sap-id [split-horizon-group group-name] [capture-sap] [create]
                — dhcp
                    — description description-string
                    — no description
                    — lease-populate [nbr-of-leases]
                    — no lease-populate
                    — [no] option
                        — action {replace | drop | keep}
                        — no action
                        — circuit-id
                        — circuit-id {ascii-tuple | vlan-ascii-tuple}
                        — circuit-id hex [0x0..0xFFFFFFFF...(64 hex nibbles)]
```

```
                        — no circuit-id
                        — remote-id
                        — remote-id hex [0x0..0xFFFFFFFF...(64 hex nibbles)]
                        — remote-id mac
                        — remote-id string string
                        — no remote-id
                        — [no] vendor-specific-option
                            — [no] client-mac-address
                            — [no] sap-id
                            — [no] service-id
                            — string text
                            — no string
                            — [no] system-id
                    — proxy-server
                        — emulated-server ip-address
                        — no emulated-server
                        — lease-time [days days] [hrs hours] [min minutes] [sec
                            seconds] [override]
                        — no lease-time
                        — [no] shutdown
                    — [no] shutdown
                    — [no] snoop
                — dhcp-user-db local-user-db-name
                — no dhcp-user-db
                — dhcp-python-policy policy-name
                — no dhcp-python-policy
                — dhcp6-user-db local-user-db-name
                — no dhcp6-user-db
            — mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [create]
                — dhcp
                    — description description-string
                    — no description
                    — [no] snoop
            — spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan}] [split-horizon-group group-
                name] [create]
                — dhcp
                    — description description-string
                    — no description
                    — [no] snoop
```

3.6.1.5.3 VPRN DHCP Commands

```
config
    — service
        — vprn
            — dhcp
            — dhcp6
                — local-dhcp-server server-name [create]
                — no local-dhcp-server server-name
                    — [no] allow-lease-query
                    — description description-string
```

- **no description**
- **failover**
  - [**no**] **ignore-mclt-on-takeover**
  - **maximum-client-lead-time** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
  - **no maximum-client-lead-time**
  - **partner-down-delay** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
  - **no partner-down-delay**
  - **peer** *ip-address* **tag** *sync-tag*
  - **no peer**
  - [**no**] **shutdown**
  - [**no**] **startup-wait-time** [**min** *minutes*] [**sec** *seconds*] [**hrs** *hours*]
- [**no**] **ignore-rapid-commit**
- [**no**] **interface-id-mapping**
- **lease-hold-time** [*lease-hold-time*]
- **no lease-hold-time**
- [**no**] **lease-hold-time-for**
  - [**no**] **internal-lease-ipsec**
  - [**no**] **solicited-release**
- **pool** *pool-name* [**create**]
- **no pool** *pool-name*
  - **delegated-prefix-length** *bits* [**minimum** *prefix-length*] [**maximum** *prefix-length*]
  - **no delegated-prefix-length**
  - **description** *description-string*
  - **no description**
  - [**no**] **exclude-prefix** *ipv6-prefix/prefix-length*
  - **failover**
    - [**no**] **ignore-mclt-on-takeover**
    - **maximum-client-lead-time** [*maximum-client-lead-time*]
    - **no maximum-client-lead-time**
    - **partner-down-delay** [*partner-down-delay*]
    - **no partner-down-delay**
    - **peer** *ip-address* **tag** *sync-tag*
    - **no peer**
    - [**no**] **shutdown**
    - [**no**] **startup-wait-time** [**min** *minutes*] [**sec** *seconds* [**hrs** *hours*]
  - **max-lease-time** [*max-lease-time*]
  - **no max-lease-time**
  - **min-lease-time** [*min-lease-time*]
  - **no min-lease-time**
  - **minimum-free** *minimum-free* [**percent**] [**event-when-depleted**]
  - **no minimum-free**
  - [**no**] **nak-non-matching-subnet**
  - **offer-time** [**min** *minutes*] [**sec** *seconds*]
  - **no offer-time**
  - **options**
    - **custom-option** *option-number* **address** *ip-address* [*ip-address*]

- — **custom-option** *option-number* **domain** *domain-string*
- — **custom-option** *option-number* **hex** *hex-string*
- — **custom-option** *option-number* **string** *ascii-string*
- — **no custom-option** *option-number*
- — **dns-server** *ip-address* [*ip-address*]
- — **no dns-server**
- — **domain-name** *domain-name*
- — **no domain-name**

— **prefix** *ipv6-addr/prefix-len* [**failover** {**local** | **remote** | **access-driven**}] [**pd**] [**wan-host**] [**create**]

— **no prefix** *ipv6-addr/prefix-len*

- — [**no**] **drain**
- — **options**
  - — **custom-option** *option-number* **address** *address* [*address*]
  - — **custom-option** *option-number* **domain** *domain-string*
  - — **custom-option** *option-number* **hex** *hex-string*
  - — **custom-option** *option-number* **string** *ascii-string*
  - — **no custom-option** *option-number*
  - — **dns-server** *ipv6-address* [*ipv6-option*]
  - — **no dns-server**
  - — **domain-name** *domain-name*
  - — **no domain-name**
  - — **lease-rebind-time** [**days** *days*] [hrs *hours*] [**min** *minutes*] [**sec** *seconds*]
  - — **no lease-rebind-time**
  - — **lease-renew-time** [**days** *days*] [hrs *hours*] [**min** *minutes*] [**sec** *seconds*]
  - — **no lease-renew-time**
  - — **lease-time** [**days** *days*] [hrs *hours*] [**min** *minutes*] [**sec** *seconds*]
  - — **no lease-time**
  - — **netbios-name-server** *ip-address* [*ip-address*]
  - — **no netbios-name-server**
  - — **netbios-node-type** *netbios-node-type*
  - — **no netbios-node-type**
- — **preferred-lifetime** [**days** *days*] [hrs *hours*] [**min** *minutes*] [**sec** *seconds*]
- — **no preferred-lifetime**
- — **rebind-timer** [*rebind-timer*]
- — **no rebind-timer**
- — **renew-timer** [*renew-timer*]
- — **no renew-timer**

— **subnet** {*ip-address*/*mask* | *ip-address netmask*} [**create**]

— **no subnet** {*ip-address*/*mask* | *ip-address netmask*}

- — [**no**] **address-range** *start-ip-address end-ip-address* [**failover** {**local** | **remote**| **access-driven**}]
- — [**no**] **drain**

— [no] exclude-addresses *start-ip-address* [*end-ip-address*]
— maximum-declined *maximum-declined*
— no maximum-declined
— minimum-free *minimum-free* [**percent**] [**event-when-depleted**]
— no minimum-free
— options
— custom-option *option-number* **address** [*ip-address*]
— custom-option *option-number* **hex** *hex-string*
— custom-option *option-number* **string** *ascii-string*
— no custom-option *option-number*
— default-router *ip-address* [*ip-address*]
— no default-router
— subnet-mask *ip-address*
— no subnet-mask
— thresholds
— [no] minimum-free **prefix-length** [*prefix-length*]
— [no depleted-event
— minimum [**percent** [*percent*]] [**number** [*number*]]
— no minimum
— valid-lifetime [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
— no valid-lifetime
— server-id **duid-en hex** *hex-string*
— server-id **duid-en string** *ascii-string*
— server-id **duid-ll**
— no server-id
— [no] shutdown
— use-gi-address [**scope** *scope*]
— no use-gi-address
— use-link-address [**scope** *scope*]
— no use-link-address
— use-pool-from-client
— use-pool-from-client **delimiter** *delimiter*
— no use-pool-from-client
— user-db *local-user-db-name*
— no user-db
— user-ident {**client-id** | **circuit-id** | **mac** | **mac-circuit-id** | **remote-id**}
— no user-ident
— [no] **interface** *ip-int-name* [**create**]
— dhcp
— description *description-string*
— no description
— gi-address *ip-address* [**src-ip-addr**]
— no gi-address
— lease-populate [*nbr-of-leases*]
— no lease-populate
— [no] option
— action {**replace** | **drop** | **keep**}

```
                    — no action
                    — circuit-id [ascii-tuple | ifindex | sap-id | vlan-ascii-tup]
                    — no circuit-id
                    — remote-id [mac | string string]
                    — no remote-id
                    — [no] vendor-specific-option
                        — [no] client-mac-address
                        — [no] pool-name
                        — [no] sap-id
                        — [no] service-id
                        — string text
                        — no string
                        — [no] system-id
                — python-policy policy-name
                — no python-policy
                — proxy-server
                    — emulated-server ip-address
                    — no emulated-server
                    — lease-time [lease-time] [override]
                    — no lease-time
                    — [no] shutdown
                — [no] relay-plain-bootp
                — relay-proxy [release-update-src-ip] [siaddr-override ip-
                      address]
                — no relay-proxy
                — server server1 [server2.]
                — no server
                — [no] shutdown
                — [no] trusted
                — [no] use-arp
```

3.6.1.5.4 IES/VPRN IPv6-DHCP6 Commands

```
configure
    — service
        — ies service-id/vprn service-id
            — subscriber-interface ip-int-name [create]
            — subscriber-interface ip-int-name [create] fwd-service service-id fwd-
                  subscriber-interface fwd-int-name
            — no subscriber-interface ip-int-name
                — [no] group-interface ip-int-name [create] [type]
                    — ipv6
                        — [no] allow-multiple-wan-addresses
                        — dhcp6
                            — filter filter-id
                            — no filter
                            — [no] option
                                — interface-id [ascii-tuple]
                                — interface-id ifindex
                                — interface-id sap-id
                                — interface-id string string
```

- **no interface-id**
- [**no**] **remote-id**

- [**no**] **override-slaac**
- **pd-managed-route** [**next-hop** {**ipv4** | **ipv6**}]
- **no pd-managed-route**
- [**no**] **proxy-server**
  - **client-applications** {[**dhcp**] [**ppp**]}
  - **no client-applications**
  - **preferred-lifetime** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
  - **preferred-lifetime infinite**
  - **no preferred-lifetime**
  - **rebind-timer** [*rebind-timer*]
  - **no rebind-timer**
  - **renew-timer** [[*renew-timer*]
  - **no renew-timer**
  - **server-id duid-en hex** *hex-string*
  - **server-id duid-en string** *ascii-string*
  - **server-id duid-ll**
  - **no server-id**
  - **no shutdown**
  - **valid-lifetime** [*valid-lifetime*]
  - **valid-lifetime infinite**
  - **no valid-lifetime**
- **python-policy** *name*
- **no python-policy**
- [**no**] **relay**
  - **client-applications** [**dhcp**] [**ppp**]
  - **no client-applications**
  - **description** *description-string*
  - **no description**
  - **link-address** *ipv6-address*
  - **no link-address**
  - **server** [*ipv6-address*]
  - **no server**
  - [**no**] **shutdown**
  - **source-address** *ipv6-address*
  - **no source-address**
- [**no**] **snooping**
- [**no**] **shutdown**
- **user-db** *local-user-db-name*
- **no user-db**
- **user-ident** *user-ident*
- **no user-ident**

3.6.1.5.5 Anti-Spoof Filtering Commands

Note:

- DHCP commands described in this section are applicable to the anti-spoofing feature. Refer to the DHCP Management chapter for information about other DHCP feature descriptions, CLI syntax, and usage.

```
config
    — service
        — ies service-id/vprn service-id
        — no ies/vprn service-id
            — interface ip-int-name [create]
            — interface ip-int-name [create] tunnel
            — no interface ip-int-name
                — sap sap-id [create]
                — no sap sap-id
                    — anti-spoof {type}
                    — no anti-spoof
                — filter
                    — description description-string
                    — no description
                    — gi-address ip-address [src-ip-addr]
                    — no gi-address
                    — lease-populate nbr-of-leases
                    — no lease-populate
                    — option
                        — action {replace | drop | keep}
                        — no action
                        — circuit-id [ascii-tuple | ifindex | sap-id | vlan-ascii-
                            tuple]
                        — no circuit-id
                        — [no] remote-id
                    — [no] relay-plain-bootp
                    — relay-proxy [release-update-src-ip] [siaddr-override ip-
                        address]
                    — no relay-proxy
                    — server server1 [server2]
                    — no server
                    — [no] shutdown
                    — [no] trusted

config
    — service
        — vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [b-vpls | i-vpls]
            [etree] [name name]
        — no vpls service-id
            — sap sap-id [split-horizon-group group-name] [create] [capture-sap] [eth-ring
                ring-index]
            — sap sap-id [split-horizon-group group-name] [create] [capture-sap] [eth-ring
                ring-index] leave-ac
```

— **sap** *sap-id* [**split-horizon-group** *group-name*] [**create**] [**capture-sap**] [**eth-ring** *ring-index*] **root-leaf-tag leaf-tag** *leaf-tag*
— **no sap** *sap-id*
— **anti-spoof** {*type*}
— **no anti-spoof**
— **dhcp**
— **description** *description-string*
— **no description**
— **lease-populate** [*nbr-of-leases*]
— **no lease-populate**
— [**no**] **option**
— **option** [*dhcp-action*]
— **no option**
— **circuit-id**
— **circuit-id** {**ascii-tuple** | **vlan-ascii-tuple**}
— **circuit-id hex** [**0x0..0xFFFFFFFF...**(**64 hex nibbles**)]
— **no circuit-id**
— [**no**] **remote-id**
— [**no**] **shutdown**
— [**no**] **snoop**
— **mesh-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**]
— **mesh-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**] **leaf-ac**
— **mesh-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**] **root-leaf-tag**
— **no mesh-sdp** *sdp-id*[:*vc-id*]
— **mesh-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**]
— **mesh-sdp** *sdp-id*[:*vc-id*] [**vc-type** {**ether** | **vlan**}] [**create**] **leaf-ac**
— **dhcp**
— **description** *description-string*
— **no description**
— [**no**] **snoop**
— **spoke-sdp** *sdp-id*[**:***vc-id*] [**vc-type** {**ether** | **vlan**}] [**split-horizon-group** *group-name*]
— **no spoke-sdp** *sdp-id*[**:***vc-id*]
— **dhcp**
— **description** *description-string*
— **no description**
— [**no**] **snoop**

3.6.1.6 Show Commands

```
— router
    — dhcp
        — local-dhcp-server server-name
            — associations
            — declined-addresses ip-address[/mask] [detail]
            — declined-addresses pool pool-name
            — failover-pool-stats [pool-name]
            — failover-server-stats
            — free-addresses ip-address[/mask]
            — free-addresses summary [subnet ip-address[/mask]
            — free-addresses pool pool-name
            — leases [detail]
            — leases ip-address[/mask] address-from-user-db [detail]
            — leases ip-address[/mask] dhcp-host dhcp-host-name [detail]
            — leases ip-address[/mask] ppp-host ppp-host-name [detail]
            — leases ip-address[/mask] [detail] [state]
            — pool-ext-stats [pool-name]
            — server-stats
            — sticky-leases [hostname]
            — subnet-ext-stats ip-address[/mask]
            — subnet-ext-stats pool pool-name
            — subnet-stats ip-address[/mask]
            — subnet-stats pool pool-name
            — summary
        — servers
        — servers all
        — statistics [interface ip-int-name | ip-address]
        — summary
— router
    — dhcp6
        — local-dhcp-server server-name
            — associations
            — failover-pool-stats [pool-name]
            — failover-server-stats
            — interface-id-mapping
            — leases [ipv6-address/prefix-length] [type] [state] [detail]
            — pool-ext-stats [pool-name]
            — pool-stats [pool-name]
            — pool-threshold-stats [pool-name] detail [format {exact | scientific}]
            — pool-threshold-stats [pool-name]
            — prefix-ext-stats ipv6-address/prefix-length
            — prefix-ext-stats pool pool-name
            — prefix-stats ipv6-address/prefix-length
            — prefix-stats pool pool-name
            — prefix-threshold-stats pool pool-name detail [format {exact |
                scientific}]
            — prefix-threshold-stats pool pool-name
            — prefix-threshold-stats ipv6-address/prefix-length detail [format
                {exact | scientific}]
            — prefix-threshold-stats ipv6-address/prefix-length
            — server-stats
```

```
                    — servers
                    — servers all
                    — statistics
                    — summary

show
    — service
        — id service-id
            — dhcp
                — lease-state [wholesaler service-id] [sap sap-id | sdp sdp-id:vc-id |
                    interface interface-name | ip-address ip-address[/mask] | chaddr
                    ieee-address | mac ieee-address| {[port port-id] [no-inter-dest-id |
                    inter-dest-id inter-dest-id]}] [session {none | ipoe}] [detail]
                — statistics [sap sap-id] | [sdp [sdp-id[:vc-id]] | interface ip-int-name]]
                — summary [interface interface-name | saps]
            — dhcp6
                — lease-state [detail] [wholesaler service-id] [session {none | ipoe |
                    ppp}] [router-advertisement-policy policy-name]
                — lease-state [detail] interface interface-name [wholesaler service-id]
                    [session {none | ipoe | ppp}] [router-advertisement-policy policy-
                    name]
                — lease-state [detail] ipv6-address ipv6-prefix[/prefix-length] [wholesaler
                    service-id] [session {none | ipoe | ppp}] [router-advertisement-
                    policy policy-name]
                — lease-state [detail] mac ieee-address [wholesaler service-id] [session
                    {none | ipoe | ppp}] [router-advertisement-policy policy-name]
                — statistics [interface interface-name]
                — summary [interface interface-name | saps]
            — virtual-subnet [sap sap-id]
            — virtual-subnet subscriber sub-ident

show
    — system
        — dhcp6
```

3.6.1.7 Tools Commands

```
tools
    — perform
        — subscriber-mgmt
            — local-user-db local-user-db-name
                — ipoe
                    — host-lookup [mac ieee-address] [remote-id remote-id] [sap-id
                        sap-id] [service-id service-id] [string vso-string]
                        [system-id system-id] [option60 hex-string] [c0ircuit-id
                        circuit-id | circuit-id-hex circuit-id-hex] [option60-hex
                        option60-hex] [remote-id-hex remote-id-hex] [derived-
                        id derived-id] [ip-prefix ip-prefix/ip-prefix-length]
                — ppp
```

```
                    — authentication user-name user-name password password
                            [circuit-id circuit-id] [circuit-id-hex circuit-id-hex]
                            [derived-id derived-id] [mac ieee-address] [remote-id
                            remote-id] [remote-id-hex remote-id-hex] [sap-id sap-id]
                            [service-name service-name]
                    — host-lookup [circuit-id circuit-id] [circuit-id-hex circuit-id-hex]
                            [derived-id derived-id] [mac ieee-address] [remote-id
                            remote-id] [remote-id-hex remote-id-hex] [sap-id sap-id]
                            [service-name service-name] [user-name user-name]
            — router
                — dhcp
                    — local-dhcp-server server-name
                        — failover
                            — abort-startup-wait
                            — force-partner-down
                        — pool pool-name
                            — create-sticky-lease host-name [mac ieee-address]
                                [circuit-id circuit-id] [client-id client-id] [requested-
                                ip-address ip-address] [circuit-id-hex circuit-id-hex-
                                string] [client-id-hex client-id-hex-string]
                            — failover
                                — abort-startup-wait
                                — force-partner-down
                        —  send-force-renew ip-address
            — router
                — dhcp6
                    — local-dhcp-server server-name
                        — failover
                            — abort-startup-wait
                            — force-partner-down
                        — pool pool-name
                            — failover
                                — abort-startup-wait
                                — force-partner-down
            — dump
                — dhcp-rx-stats [clear]
```

3.6.1.8 Clear Commands

```
clear
    — router
        — dhcp
            — local-dhcp-server server-name
                — declined-addresses ip-address[/mask]
                — declined-addresses pool pool-name
                — failover-pool-stats [pool-name]
                — leases ip-address[/mask] [state]
                — leases all [state]
                — pool-ext-stats [pool-name]
                — server-stats
                — sticky-leases hostname hostname
```

```
                        — sticky-leases hostname-prefix hostname-prefix
                        — subnet-ext-stats ip-address[/mask]
                        — subnet-ext-stats pool pool-name
                    — statistics [ip-int-name | ip-address]
                — dhcp6
                    — local-dhcp-server server-name
                        — failover-pool-stats [pool-name]
                        — leases [ipv6-address/prefix-length] [type] [state]
                        — leases all [type] [state]
                        — pool-ext-stats [pool-name]
                        — pool-threshold-stats [pool-name]
                        — pool-ext-stats ipv6-address/prefix-length
                        — prefix-ext-stats pool pool-name
                        — prefix-threshold-stats ipv6-address/prefix-length
                        — prefix-threshold-stats pool pool-name
                        — server-stats
                    — statistics
            — service
                — id service-id
                    — dhcp
                        — lease-state all [no-dhcp-release]
                        — lease-state [port port-id] [inter-dest-id intermediate-destination-id]
                          [no-dhcp-release]
                        — lease-state [port port-id] no-inter-dest-id [no-dhcp-release]
                        — lease-state ip-address ip-address[/mask] [no-dhcp-release]
                        — lease-state mac ieee-address [no-dhcp-release]
                        — lease-state sap sap-id [no-dhcp-release]
                        — lease-state sdp sdp-id:vc-id [no-dhcp-release]
                        — statistics [sap sap-id | sdp [sdp-id[:vc-id] | interface ip-int-name | ip-
                          address]
                    — dhcp6
                        — lease-state all [no-dhcp-release]
                        — lease-state ipv6-address ipv6-prefix[/prefix-length] [no-dhcp-release]
                        — lease-state mac ieee-address [no-dhcp-release]
                        — lease-state sap sap-id [no-dhcp-release]
                        — statistics [interface ip-int-name | ipv6-address]
```

3.6.1.9 Debug Commands

```
debug
    — router router-instance
        — ip
            — [no] dhcp [interface ip-int-name]
            — [no] dhcp mac ieee-address
            — [no] dhcp sap sap-id
                — detail-level {low | medium | high}
                — no detail-level
                — mode {dropped-only | ingr-and-dropped | egr-ingr-and-dropped}
                — no mode
            — dhcp6 [interface ip-int-name]
            — no dhcp6
```

- **detail-level** {**low** | **medium** | **high**}
- **no detail-level**
- **mode** {**dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}
- **no mode**

— [**no**] **local-dhcp-server** *server-name* [**lease-address** *ip-prefix*[*/prefix-length*]]
— [**no**] **local-dhcp-server** *server-name* [**mac** *ieee-address*]
— [**no**] **local-dhcp-server** *server-name* **link-local-address** *ipv6z-address*

- **detail-level** {**low** | **medium** | **high**}
- **no detail-level**
- **mode** {**dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}
- **no mode**

3.6.2 Command Descriptions

3.6.2.1 Triple Play DHCP Configuration Commands

3.6.2.1.1 Global Commands

description

Syntax **description** *description-string*
no description

Context config>filter>dhcp-filter
config>router>dhcp>server
config>router>dhcp>server>pool
config>service>ies>if>dhcp
config>service>ies>if>ipv6>dhcp6-relay
config>service>ies>sub-if>grp-if>ppp
config>service>ies>sub-if>ipv6>dhcp6>relay
config>service>vpls>sap>dhcp
config>service>vpls>sap>ipoe-session
config>service>vprn>if>dhcp
config>service>vprn>sub-if>ipv6>dhcp6>relay
config>subscr-mgmt>loc-user-db

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Parameters *description-string* — The description character string. Allowed values are any string, up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

shutdown

Syntax [**no**] **shutdown**

Context config>service>ies>if>dhcp
config>service>vpls>sap>dhcp
config>service>vpls>sap>dhcp>option>vendor
config>service>vpls>sap>ipoe-session
config>service>vprn>if>dhcp
config>service>vprn>if>dhcp>proxy-server
config>subscr-mgmt>loc-user-db
config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host
config>router>dhcp6>server>failover
config>router>dhcp>server>failover

Description This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command places the entity into an administratively enabled state.

3.6.2.1.2 System DHCP Commands

dhcp6

Syntax **dhcp6**

Context config>system

Description This command enables the context to configure DHCP6 parameters.

adv-noaddrs-global

Syntax **adv-noaddrs-global** [**esm-proxy**] [**esm-relay**] [**relay**] [**server**]
no adv-noaddrs-global

Context config>system>dhcp6

Description This command configures the different DHCPv6 applications to send the NoAddrsAvail Status-Code in DHCPv6 Advertise messages at the global DHCP message level.

By default, all applications send the NoAddrsAvail Status-Code in DHCPv6 Advertise messages at the IA_NA Option level.

Different applications for which NoAddrsAvail Status-Code in DHCPv6 Advertise messages can be configured at the global DHCP message level.

The only valid combination in current SR OS is **adv-noaddrs-global esm-relay server**.

The **no** form of this command reverts to the default.

Default no adv-noaddrs-global. All applications send the NoAddrsAvail Status-Code in DHCPv6 Advertise messages at the IA_NA Option level.

Parameters **esm-proxy** — Specifies the DHCPv6 proxy server on subscriber group-interfaces. Not supported in current SR OS.

esm-relay — Specifies the DHCPv6 relay on subscriber group-interfaces. Must be enabled together with the DHCPv6 server (server) application.

relay — Specifies the DHCPv6 relay on regular IES or VPRN interfaces. Not supported in current SR OS.

server — Specifies the DHCPv6 server. Must be enabled together with the DHCPv6 relay on subscriber interfaces (esm-relay) application.

3.6.2.1.3 DHCP Filter Commands

dhcp-filter

Syntax **dhcp-filter** *filter-id* [**create**]
no dhcp-filter

Context config>filter

Description This command creates a configuration context for the specified DHCP filter policy if it does not exist and enables the context to configure the specified DHCP filter policy.

Parameters *filter-id* — Specifies the DHCP filter policy ID expressed as a decimal integer.

Values 1 to 65535

create — Keyword required to create the configuration context.

dhcp6-filter

Syntax **dhcp6-filter** *filter-id* [**create**]
no dhcp6-filter

Context config>filter

Description This command creates a configuration context for the specified DHCPv6 filter policy if it does not exist and enables the context to configure the specified DHCPv6 filter policy.

The **no** form of this command reverts to the default.

Parameters *filter-id* — Specifies the DHCPv6 filter policy ID expressed as a decimal integer.

Values 1 to 65535

create — Keyword required to create the configuration context.

default-action

Syntax **default-action bypass-host-creation**
default-action drop
no default-action

Context config>filter>dhcp-filter

Description This command specifies the default action when no entries match.

The **no** form of this command reverts to the default.

Parameters **bypass-host-creation** — Specifies to bypass ESM host creation options.

drop — Specifies to drop and not process the DHCP message.

default-action

Syntax **default-action bypass-host-creation** [**na**] [**pd**]
default-action drop
no default-action

Context config>filter>dhcp6-filter

Description This command specifies the default action when no entries match.

The **no** form of this command reverts to the default.

Parameters **bypass-host-creation** — Specifies to bypass ESM host creation options.

Values **na** — Bypasses the DHCP NA hosts creation.

pd — Bypasses the DHCP PD hosts creation.

drop — Specifies to drop and not process the DHCP6 message.

entry

Syntax **entry** *entry-id* [**create**]
no entry *entry-id*

Context config>filter>dhcp-filter
config>filter>dhcp6-filter

Description This command configures DHCP filter entries.

The **no** form of this command removes the entry from the configuration.

Parameters *entry-id* — Specifies the entry ID.

Values 1 to 65535

create — This keyword is required when first creating the DHCP filter entry. Once the context is created, it is possible to navigate into the context without the **create** keyword.

action

Syntax **action bypass-host-creation**
action drop
no action

Context config>filter>dhcp-filter>entry

Description This command specifies the action to take on DHCP host creation when the filter entry matches.

The **no** form of this command reverts to the default wherein the host creation proceeds as normal.

Parameters **bypass-host-creation** — Specifies that the host creation is bypassed.

drop — Specifies that the DHCP message is dropped.

action

Syntax **action bypass-host-creation** [**na**] [**pd**]
action drop
no action

Context config>filter>dhcp6-filter>entry

Description This command specifies the action to take on DHCP6 host creation when the filter entry matches.

The **no** form of this command reverts to the default wherein the host creation proceeds as normal.

Parameters **bypass-host-creation** — Specifies that the host creation is bypassed.

Values **na** — Bypasses the DHCP6 NA hosts creation.
pd — Bypasses the DHCP6 PD hosts creation.

drop — Specifies that the DHCP6 message is dropped.

action

Syntax **action** {**replace** | **drop** | **keep**}
no action

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option

Description This command configures the Relay Agent Information Option (Option 82) processing.

The **no** form of this command returns the system to the default value.

Default The default is to keep the existing information intact.

Parameters **replace** — Specifies to overwrite any existing Option 82 field. In the upstream direction (from the user), the Option 82 field from the router is inserted in the packet. In the downstream direction (toward the user) the Option 82 field is stripped (in accordance with RFC 3046).

drop — Specifies to drop the DHCP packet if an Option 82 field is present, and a counter is incremented.

keep — Specifies to keep existing information in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is forwarded towards the client.

option

Syntax **option** *dhcp-option-number* {**present** | **absent**}
option *dhcp-option-number* **match hex** *hex-string* [**exact**] [**invert-match**]
option *dhcp-option-number* **match string** *ascii-string* [**exact**] [**invert-match**]
no option

Context config>filter>dhcp-filter>entry

Description This command configures match criteria for the DHCP filter policy entry.

The **no** form of this command reverts to the default

Parameters *dhcp-option-number* — Specifies the DHCP option number.

Values 0 to 255

present — Specifies that the related DHCP option must be present.

absent — Specifies that the related DHCP option must be absent.

match hex *hex-string* — The option must partially match a specified hex string.

Values 0x0 to 0xFFFFFFFF (up to 254 hex nibbles)

match string *ascii-string* — The option must partially match a specified ASCII string, up to 127 characters.

exact — This option requires an exact match of a hex or ASCII string.

invert-match — Requires the option not to partially match.

option

Syntax **option** *dhcp6-option-number* {**present** | **absent**}
option *dhcp6-option-number* **match hex** *hex-string* [**exact**] [**invert-match**]
option *dhcp6-option-number* **match string** *ascii-string* [**exact**] [**invert-match**]
no option

Context config>filter>dhcp6-filter>entry

Description This command configures match criteria for the DHCP6 filter policy entry.

The **no** form of this command reverts to the default.

Parameters *dhcp-option-number* — Specifies the DHCP6 option number.

Values 0 to 255

present — Specifies that the related DHCP6 option must be present.

absent — Specifies that the related DHCP6 option must be absent.

match hex *hex-string* — The option must (partially) match a specified hex string.

Values 0x0 to 0xFFFFFFFF (up to 254 hex nibbles)

match string *ascii-string* — The option must partially match a specified ASCII string, up to 127 characters.

exact — This option requires an exact match of a hex or ASCII string.

invert-match — Requires the option not to partially match.

3.6.2.1.4 Local User Database Commands

subscriber-mgmt

Syntax **subscriber-mgmt**

Context config>subscr-mgmt

Description This command enables the context to configure subscriber management parameters.

local-user-db

Syntax **local-user-db** *local-user-db-name* [**create**]
no local-user-db *local-user-db-name*

Context config>subscr-mgmt

Description This command enables the context to configure a local user database.

The **no** form of this command reverts to the default.

Parameters *local-user-db-name* — Specifies the name of a local user database.

ipoe

Syntax **ipoe**

Context config>subscr-mgmt>loc-user-db

Description This command enables the context to configure IPoE host parameters.

host

Syntax **host** *host-name* [**create**]
no host *host-name*

Context config>subscr-mgmt>loc-user-db>ipoe
config>subscr-mgmt>loc-user-db>ppp

Description This command defines a DHCP or PPP subscriber.

The **no** form of this command removes the host name from the configuration.

Parameters *host-name* — Specifies the host name.

create — Keyword used to create the host name. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

acct-policy

Syntax **acct-policy** *acct-policy-name* [**duplicate** *acct-policy-name*]
no acct-policy

Context config>subscr-mgmt>loc-user-db>ppp>host
config>subscr-mgmt>loc-user-db>ipoe>host

Description This command specifies the accounting policy used for sending an Accounting Stop message to report RADIUS authentication failures of PPPoE sessions. A duplicate policy can be specified if a copy of the Accounting Stop message must be sent to another destination.

Reporting RADIUS authentication failures with an Accounting Stop message must be enabled in the RADIUS authentication policy (“send-acct-stop-on-fail”).

A duplicate RADIUS accounting policy can be specified if the accounting stop resulting from a RADIUS authentication failure must also be sent to a second RADIUS destination.

The **no** form of this command reverts to the default.

Parameters *acct-policy-name* — Specifies the name of a RADIUS accounting policy, up to 32 characters.

address

Syntax **address gi-address** [**scope** *scope*]
address *ip-address*[*/prefix-length*]
address pool *pool-name* [**secondary-pool** *sec-pool-name*] [**delimiter** *delimiter*]
address use-pool-from-client [**delimiter** *delimiter*]
no address

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command configures how the IP address is defined for this host.

When the user database is used from a local DHCP server, then this command defines how to define the IP address the server will “offer” to the DHCP-client.

When the user-db is used for PPPoE authentication, the **gi-address** parameter cannot be used. A fixed IP address will then cause PPPoE to use this IP address. If no IP address is specified, the PPPoE looks for IP address by other means (DHCP). If a pool name is given, this pool will be sent in the DHCP request so that it can be used in by the DHCP server to determine which address to give to the host.

The **no** form of this command causes no IP address to be assigned to this host. In a user database referred to from a local DHCP server, creating a host without address information will cause the matching client never to get an IP address.

The **no** form of this command reverts to the default.

Parameters **gi-address** — When specified, the gi-address of the DHCP message is taken to look for a subnet in the local DHCP server. The first available free address of the subnet is taken and "offered" to the host. When **local-user-db** is used for PPPoE authentication, this has the same result as **no address**.

ip-address — Specifies the fixed IP address to use for this host.

pool-name/sec-pool-name — Specifies the primary (and secondary) pool (in the local DHCP server) to use to look for an available address. The first available IP address from any subnet in the pool will be used. When the local user database is used for PPPoE authentication, this causes the specified pool name to be sent to the DHCP server in a vendor-specific sub-option under Option 82.

use-pool-from-client — Use the pool-name in the Option 82 vendor-specific sub-option.

delimiter — A single ASCII character specifies the delimiter of separating primary and secondary pool names in option82 VSO.

auth-domain-name

Syntax **auth-domain-name** *domain-name*
no auth-domain-name

Context config>subscr-mgmt>loc-user-db>ipoe>host

Description This command sets the domain name which can be appended to user-name in RADIUS-authentication-request message for the given host.

The **no** form of this command removes the domain name from the host configuration.

Parameters *domain-name* — Specifies the domain name, up to 32 characters, to be appended to user-name in RADIUS-authentication-request message for the given host.

auth-policy

Syntax **auth-policy** *policy-name*
no auth-policy

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command configures the authentication policy of this host and PPPoE hosts. This authentication policy is only used if no authentication policy is defined at the interface level. For DHCP hosts, the host entry should not contain any other information needed for setup of the host (IP address, ESM strings, and so on.). For PPPoE hosts, the authentication policy configured here must have its PPPoE authentication method set to **pap-chap**, otherwise the request is dropped.

The **no** form of this command reverts to the default.

Parameters *policy-name* — Specifies the authentication policy name, up to 32 characters.

diameter-application-policy

Syntax **diameter-application-policy** *policy-name*
no diameter-application-policy

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command configures the Diameter application policy.

The **no** form of this command reverts to the default.

Parameters *policy-name* — Specifies the Diameter application policy name.

diameter-auth-policy

Syntax **diameter-auth-policy** *name*
no diameter-auth-policy

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command is used to configure the Diameter NASREQ application policy to use for authentication.

The **no** form of this command reverts to the default.

Parameters *name* — Specifies the name of the Diameter NASREQ application policy to use for authentication.

gi-address

Syntax **gi-address** *ip-address*
no gi-address

Context config>subscr-mgmt>loc-user-db>ipoe>host

Description This command allows selection of gi-addresses based on the host entry in LUDB.

The gi-address must be a valid address (associated with an interface) within the routing context that received the DHCP message on the access side.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the IPv4 gi-address.

host-identification

Syntax **host-identification**

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command enables the context to configure host identification parameters.

circuit-id

Syntax **circuit-id string** *ascii-string*
circuit-id hex *hex-string*
no circuit-id

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident
config>subscr-mgmt>loc-user-db>ppp>host>host-ident

Description This command specifies the circuit-id to match.

The **no** form of this command reverts to the default.

Parameters *ascii-string* — Specifies the circuit ID from the Option 82.

hex-string — Specifies the circuit ID in hexadecimal format from the Option 82.

Values 0x0 to 0xFFFFFFFF (maximum 254 hex nibbles)

derived-id

Syntax **derived-id** *derived-id-string*
no derived-id

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident
config>subscr-mgmt>loc-user-db>ppp>host>host-ident

Description This command configures an ASCII string that uniquely identifies a host, and is derived by a Python script from packet content available during a DHCP transaction.

The **no** form of this command reverts to the default.

Parameters *derived-id-string* — Specifies the host ID to be derived by a python script from DHCP packets during a DHCP transaction, up to 255 characters.

encap-tag-range

Syntax **encap-tag-range start-tag** *start-tag* **end-tag** *end-tag*
no encap-tag-range

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident
config>subscr-mgmt>loc-user-db>ppp>host>host-ident

Description This command specifies a range of encapsulation tag as the host identifications. The encapsulation tag is dot1q or qinq on Ethernet port; VPI/VCI on ATM port.

For dot1q, the start/end-tag is single number, range from 0-4094; for QinQ, the start/end-tag format is x.y, x or y could be “*”, which means ignore inner or outer tag;

For ATM the start/end-tag format is vpi/vci, vpi or vci could be “*”, which means ignore VPI or VCI.

Note: This command will only be used when **encap-tag-range** is configured as one of the **match-list** parameters.

The **no** form of this command reverts to the default.

Parameters **start-tag** *start-tag* — Specifies the value of the start label in the range of SAP’s allowed on this host.

Values

start-tag	dot1q	qtag1	
	qinq	(qtag1.qtag2 | qtag1.* | *.qtag2)	
	atm	(vpi/vci | vpi/* | */vci)	
		qtag1	[0 to 4094]
		qtag2	[0 to 4094]
		vpi	[0 to 4095] (NNI)
			[0 to 255] (UNI)
		vci	[1 to 65535]

end-tag *end-tag* — Specifies the value of the end label in the range of SAP's allowed on this host.

Values

end-tag	dot1q	qtag1
	qinq	(qtag1.qtag2 | qtag1.* | *.qtag2)
	atm	(vpi/vci | vpi/* | */vci)
		qtag1 [0 to 4094]
		qtag2 [0 to 4094]
		vpi [0 to 4095] (NNI)
		[0 to 255] (UNI)
		vci [1 to 65535]

ip-prefix

Syntax **ip-prefix** *ip-prefix/ip-prefix-length*
no ip-prefix

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Description This command specifies the source IPv4/IPv6 address/prefix of the data trigger packet as the host identification.

The **no** form of this command reverts to the default.

Parameters *ip-prefix/ip-prefix-length* — Specifies the IPv4 address, IPv6 address, or IPv6 prefix.

Values

ipv4-prefix	a.b.c.d
ipv4-prefix-length	[0 to 32]
ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x to [0 to FFFF]H
	d to [0 to 255]D
ipv6-prefix-length	[0 to 128]

mac

Syntax **mac** *ieee-address*
no mac

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident
config>subscr-mgmt>loc-user-db>ppp>host>host-ident

Description This command specifies the MAC address to match.

The **no** form of this command reverts to the default.

Parameters *ieee-address* — Specifies the 48-bit MAC address in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers.

option60

Syntax **option60 hex** *hex-string*
option60 string *ascii-string*
no option60

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Description This command specifies the Vendor-Identifying Vendor Option to match. Option 60 is encoded as Type-Length-Value (TLV). The *hex-string* portion of Option 60 in the received DHCP request is used for matching. Only the first 32 bytes can be defined here. If Option 60 from the message is longer, those bytes are ignored.

The **no** form of this command reverts to the default.

Parameters *hex-string* — Specifies the hexadecimal format for this option.

Values 0x0 to 0xFFFFFFFF(maximum 64 hex nibbles)

ascii-string — Specifies the string format for this option, up to 32 characters.

remote-id

Syntax **remote-id hex** *hex-string*
remote-id string *ascii-string*
no remote-id

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident
config>subscr-mgmt>loc-user-db>ppp>host>host-ident

Description This command specifies the remote ID of this host.

The **no** form of this command returns to the default.

Parameters *hex-string* — Specifies the hexadecimal format for the remote ID.

Values 0x0 to 0xFFFFFFFF (maximum 254 hex nibbles)

ascii-string — Specifies the string format for the remote ID, up to 255 characters.

sap-id

Syntax **sap-id** *sap-id*

no sap-id

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Description This command specifies the SAP ID from the Nokia vendor-specific sub-option in Option 82 to match.

The **no** form of this command returns to the default.

Parameters *sap-id* — Specifies a SAP ID, up to 255 characters.

service-id

Syntax **service-id** *service-id*
no service-id

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Description This command specifies an existing service ID from the Nokia vendor-specific sub-option in Option 82 to match.

The **no** form of this command returns to the default.

Parameters *service-id* — Specifies an existing service ID or service name.

Values *service-id* — 1 to 2147483647
service-name — up to 64 characters

string

Syntax **string** *string*
no string

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Description This command specifies the string from the Nokia vendor-specific sub-option (VSO) in Option 82 to match.

The **no** form of this command returns to the default.

Parameters *string* — Specifies the VSO string of this host, up to 255 characters.

system-id

Syntax **system-id** *system-id*
no system-id

Context config>subscr-mgmt>loc-user-db>ipoe>host>host-ident

Description This command specifies the system ID from the Nokia vendor-specific sub-option in Option 82 to match.

The **no** form of this command returns to the default.

Parameters *system-id* — Specifies the system ID, up to 255 characters.

identification-strings

Syntax **identification-strings** *option-number* [**create**]
no identification-strings

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command specifies identification strings for the subscriber. This is useful when the server is centralized with Enhanced Subscriber Management (ESM) in a lower level in the network. These strings will be parsed by a downstream Python script or they can be used literally if the **strings-from-option** option in the **config>subscr-mgmt>sub-ident-policy** context is set to this option number. In this case, the option number may be set to any allowed number (between 224 and 254 is suggested, as these are not dedicated to specific purposes). If the option number is not given, a default value of 254 is used. For PPPoE only, if the local user database is attached to the PPPoE node under the group interface and not to a local DHCP server, the strings will be used internally so the option number is not used.

The **no** form of this command returns to the default.

Parameters *option-number* — Specifies identification strings for the subscriber.

Values 1 to 254

ancp-string

Syntax **ancp-string** *ancp-string*
no ancp-string

Context config>subscr-mgmt>loc-user-db>ipoe>host>ident-strings
config>subscr-mgmt>loc-user-db>ppp>host>ident-strings

Description This command specifies the ANCP string which is encoded in the identification strings.

The **no** form of this command returns to the default.

Parameters *ancp-string* — Specifies the ANCP string, up to 63 characters.

app-profile-string

Syntax **app-profile-string** *app-profile-string*
no app-profile-string

Context config>subscr-mgmt>loc-user-db>ipoe>host>ident-strings
config>subscr-mgmt>loc-user-db>ppp>host>ident-strings

Description This command specifies the application profile string which is encoded in the identification strings.

The **no** form of this command returns to the default.

Parameters *app-profile-string* — Specifies the application profile string, up to 16 characters.

category-map-name

Syntax **category-map-name** *category-map-name* [**create**]
[no] category-map-name *category-map-name*

Context config>subscr-mgmt>loc-user-db>ipoe>host>ident-strings
config>subscr-mgmt>loc-user-db>ppp>host>ident-strings

Description This command specifies the category map name.

The **no** form of this command removes the category map name from the configuration.

Parameters *category-map-name* — Specifies an existing category map name, up to 32 characters.

inter-dest-id

Syntax **inter-dest-id** *intermediate-destination-id*
no inter-dest-id

Context config>subscr-mgmt>loc-user-db>ipoe>host>ident-strings
config>subscr-mgmt>loc-user-db>ppp>host>ident-strings

Description This command specifies the intermediate destination identifier which is encoded in the identification strings.

The **no** form of this command returns to the default.

Parameters *intermediate-destination-id* — Specifies the intermediate destination identifier, up to 32 characters.

sla-profile-string

Syntax **sla-profile-string** *sla-profile-string*
no sla-profile-string

Context config>subscr-mgmt>loc-user-db>ipoe>host>ident-strings
config>subscr-mgmt>loc-user-db>ppp>host>ident-strings

Description This command specifies the SLA profile string which is encoded in the identification strings.

The **no** form of this command returns to the default.

Parameters *sla-profile-string* — Specifies the SLA profile string, up to 16 characters.

sub-profile-string

Syntax **sub-profile-string** *sub-profile-string*
no sub-profile-string

Context config>subscr-mgmt>loc-user-db>ipoe>host>ident-strings
config>subscr-mgmt>loc-user-db>ppp>host>ident-strings

Description This command specifies the subscriber profile string which is encoded in the identification strings.

The **no** form of this command returns to the default.

Parameters *sub-profile-string* — Specifies the subscriber profile string, up to 16 characters.

ancp-string

Syntax **ancp-string** *ancp-string*
no ancp-string

Context config>subscr-mgmt>loc-user-db>ipoe>host>ident-strings
config>subscr-mgmt>loc-user-db>ppp>host>ident-strings

Description This command specifies the ANCP string which is encoded in the identification strings.

The **no** form of this command returns to the default.

Parameters *ancp-string* — Specifies the ANCP string, up to 63 characters.

app-profile-string

Syntax **app-profile-string** *app-profile-string*

no app-profile-string

Context config>subscr-mgmt>loc-user-db>ipoe>host>ident-strings
config>subscr-mgmt>loc-user-db>ppp>host>ident-strings

Description This command specifies the application profile string which is encoded in the identification strings.

The **no** form of this command returns to the default.

Parameters *app-profile-string* — Specifies the application profile string, up to 16 characters.

category-map-name

Syntax **category-map-name** *category-map-name* [**create**]
[no] category-map-name *category-map-name*

Context config>subscr-mgmt>loc-user-db>ipoe>host>ident-strings
config>subscr-mgmt>loc-user-db>ppp>host>ident-strings

Description This command specifies the category map name.

The **no** form of this command removes the category map name from the configuration.

Parameters *category-map-name* — Specifies an existing category map name, up to 32 characters.

inter-dest-id

Syntax **inter-dest-id** *intermediate-destination-id*
no inter-dest-id

Context config>subscr-mgmt>loc-user-db>ipoe>host>ident-strings
config>subscr-mgmt>loc-user-db>ppp>host>ident-strings

Description This command specifies the intermediate destination identifier which is encoded in the identification strings.

The **no** form of this command returns to the default.

Parameters *intermediate-destination-id* — Specifies the intermediate destination identifier, up to 32 characters.

sla-profile-string

Syntax **sla-profile-string** *sla-profile-string*
no sla-profile-string

Context config>subscr-mgmt>loc-user-db>ipoe>host>ident-strings
config>subscr-mgmt>loc-user-db>ppp>host>ident-strings

Description This command specifies the SLA profile string which is encoded in the identification strings.

The **no** form of this command returns to the default.

Parameters *sla-profile-string* — Specifies the SLA profile string, up to 16 characters.

spi-sharing-group-id

Syntax **spi-sharing-group-id** *group-id*
no spi-sharing-group-id

Context config>subscr-mgmt>loc-user-db>ipoe>host>ident-strings
config>subscr-mgmt>loc-user-db>ppp>host>ident-strings

Description This command configures the SLA Profile Instance (SPI) sharing group identifier for an IPoE or PPPoE session. It overrides the default SPI sharing method (**def-instance-sharing**) configured in the SLA profile.

When an spi-sharing-group-id is configured, the IPoE or PPPoE session shares the SLA Profile Instance with other IPoE or PPPoE sessions from the same subscriber that: have the same SLA profile associated, are active on the same SAP, and have the same group identifier.

Configuring an **spi-sharing-group-id** *group-id* for an IPoE host, when the IPoE session is disabled on the group interface, results in a setup failure.

The **no** form of this command returns the SPI sharing group identifier to its default.

Parameters *group-id* — Specifies the SPI group identifier.

Values 0 to 65535

sub-profile-string

Syntax **sub-profile-string** *sub-profile-string*
no sub-profile-string

Context config>subscr-mgmt>loc-user-db>ipoe>host>ident-strings
config>subscr-mgmt>loc-user-db>ppp>host>ident-strings

Description This command specifies the subscriber profile string which is encoded in the identification strings.

The **no** form of this command returns to the default.

Parameters *sub-profile-string* — Specifies the subscriber profile string, up to 16 characters.

subscriber-id

Syntax **subscriber-id** *sub-ident-string*
no subscriber-id

Context config>subscr-mgmt>loc-user-db>ipoe>host>ident-strings
config>subscr-mgmt>loc-user-db>ppp>host>ident-strings

Description This command specifies the subscriber ID which is encoded in the identification strings.

The **no** form of this command returns to the default.

Parameters *sub-ident-string* — Specifies the subscriber ID string, up to 32 characters.

service-name

Syntax **service-name** *service-name*
no service-name

Context config>subscr-mgmt>loc-user-db>ppp>host>host-ident

Description This command specifies the service name tag in PADI and/or PADR packets to match for PPPoE hosts.

The **no** form of this command returns to the default.

Parameters *service-name* — Specifies a PPPoE service name, up to 255 characters.

snooping

Syntax [**no**] **snooping**

Context config>service>ies>sub-if>grp-if>ipv6>dhcp6
config>service>vprn>sub-if>grp-if>ipv6>dhcp6

Description This command enables the group-interface to snoop DHCPv6 relay messages exchange between the subscriber host and the DHCPv6 server. A successful DHCPv6 address assignment will trigger ESM DHCPv6 host creation and a release of the lease will trigger host deletion. This feature is for ESMv6 applications where a Layer 3 aggregation network is upstream from the BNG.

The **no** form of this command reverts to the default.

user-db

Syntax **user-db** *local-user-db-name*

no user-db

Context config>service>ies>sub-if>grp-if>dhcp
config>service>ies>sub-if>grp-if>dhcp6
config>service>ies>sub-if>grp-if>ipv6>dhcp6
config>service>ies>sub-if>grp-if>ipv6>dhcp6

Description This command assigns a local user database.

The **no** form of this command reverts to the default.

Parameters *local-user-db-name* — Specifies the name of a local user database, up to 32 characters.

user-ident

Syntax **user-ident** *user-ident*
no user-ident

Context config>service>ies>sub-if>grp-if>ipv6>dhcp6
config>service>vprn>sub-if>grp-if>ipv6>dhcp6

Description This feature is only applicable when DHCPv6-snooping is enabled. The Ethernet header MAC address on DHCPv6 is used as the default key host identification. This command allows addition the keys for identifying the DHCPv6 host. The interface-id can be included in addition to the MAC key to further differentiate each DHCPv6 host.

The **no** form of this command reverts to the default.

Default user-ident mac

Parameters *user-ident* — Specifies the DHCP6 user-identification for this interface.

Values **mac** — Specifies to use only the Ethernet MAC of the DHCPv6 message to identify the host.
mac-interface-id — Specifies to additionally use the interface-id to identify the DHCPv6 host.

mask

Syntax **mask type** *ipoe-match-type* {[**prefix-string** *prefix-string* | **prefix-length** *prefix-length*] [**suffix-string** *suffix-string* | **suffix-length** *suffix-length*]}
no mask type *ipoe-match-type*

Context config>subscr-mgmt>loc-user-db>ipoe

Description This command configures the IPoE mask.

The **no** form of this command removes the parameters from the mask configuration.

Parameters *match-type* — Specifies up to four matching types to identify a host.

Values circuit-id, option60, remote-id, sap-id, string, system-id

prefix-string — Specifies a substring that is stripped of the start of the incoming circuit ID before it is matched against the value configured in the circuit ID. This string can only contain printable ASCII characters. The "*" character is a wildcard that matches any substring. If a "\" character is masked, use the escape key so it becomes "\\".

Values 127 characters maximum, *' is a wildcard

prefix-length — Specifies the number of characters to remove from the start of the incoming circuit ID before it is matched against the value configured in the circuit ID.

Values 1 to 127

suffix-string — Specifies a substring that is stripped of the end of the incoming circuit ID before it is matched against the value configured in the circuit ID.

This string can only contain printable ASCII characters. The "*" character is a wildcard that matches any substring. If a "\" character is masked, use the escape key so it becomes "\\".

Values 127 characters maximum

suffix-length — Specifies the number of characters to remove from the end of the incoming circuit ID before it is matched against the value configured in the circuit ID.

Values 1 to 127

mask

Syntax **mask type** *ppp-match-type* {[**prefix-string** *prefix-string* | **prefix-length** *prefix-length*] [**suffix-string** *suffix-string* | **suffix-length** *suffix-length*]}
no mask type *ppp-match-type*

Context config>subscr-mgmt>loc-user-db>ppp

Description This command configures the mask.

The **no** form of this command removes the parameters from the mask configuration.

Parameters *ppp-match-type* — Specifies up to four matching types to identify a host.

Values circuit-id, remote-id, sap-id, service-name, username

prefix-string — Specifies a substring that is stripped of the start of the incoming circuit ID before it is matched against the value configured in the circuit ID. This string can only contain printable ASCII characters. The "*" character is a wildcard that matches any substring. If a "\" character is masked, use the escape key so it becomes "\\".

Values 127 characters maximum, *' is a wildcard

prefix-length — Specifies the number of characters to remove from the start of the incoming circuit ID before it is matched against the value configured in the circuit ID.

Values 1 to 127

suffix-string — Specifies a substring that is stripped of the end of the incoming circuit ID before it is matched against the value configured in circuit ID.

This string can only contain printable ASCII characters. The "*" character is a wildcard that matches any substring. If a "\" character is masked, use the escape key so it becomes "\\".

Values 127 characters maximum

suffix-length — Specifies the number of characters to remove from the end of the incoming circuit ID before it is matched against the value configured in the circuit ID.

Values 1 to 127

server

Syntax **server** *ip-address*
no server

Context config>subscr-mgmt>loc-user-db>ipoe>host

Description This command configures the IP address of the DHCP server to relay to.

The configured DHCP server IP address must reference one of the addresses configured under the DHCP CLI context of an IES/VPRN subscriber or group interface.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the IP address of the DHCP server.

username

Syntax **username** *user-name*
username *user-name* **no-domain**
username *user-name* **domain-only**
no username

Context config>subscr-mgmt>loc-user-db>ppp>host>host-ident

Description This command specifies how the username is specified.

The **no** form of this command returns to the default.

Parameters *username* — Specifies the user name of this host.

no-domain — No username is specified.

domain-only — Only the domain part of the username is specified, for example, nokia.com.

interface

Syntax **interface** *ip-int-name* **service-id** *service-id*
no interface

Context config>subscr-mgmt>loc-user-db>ipoe>host

Description This command specifies the interface where IPoE sessions are terminated.

The **no** version of this command disables the parameter.

Parameters *ip-int-name* — Specifies the name of the group interface.

service-id — Specifies the service ID or name where the group interface resides.

Values *service-id* — 1 to 2147483647
service-name — up to 64 characters

ipv6-address

Syntax **ipv6-address** *ipv6-address*
no ipv6-address

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command configures static DHCPv6 IA-NA address for the host. This address is delegated to the client as /128 via DHCPv6 proxy function within the router. This IP address must not be part of any DHCP pool within internal DHCP server.

The **no** form of this command removes the IPv6 address from the host configuration.

Parameters *ipv6-address* — Specifies the IPv6 address.

Values

ipv6-address:	ipv6-prefix x:x:x:x:x:x:x:x	(eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d	
	x [0 to FFFF]H	
	d [0 to 255]D	

ipv6-delegated-prefix

Syntax **ipv6-delegated-prefix** *ipv6-prefix/prefix-length*

no ipv6-delegated-prefix

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command configures static DHCPv6 IA-PD prefix for the host. This prefix can be further delegated by the host itself to its clients. The prefix length is restricted to 48 to 64 bits. This prefix must not be part of any DHCP pool within internal DHCP server.

Parameters *ipv6-address* — Specifies the IPv6 address.

Values

ipv6-address:	ipv6-prefix x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x [0 to FFFF]H
	d [0 to 255]D
prefix-length	[48 to 64]

ipv6-delegated-prefix-length

Syntax **ipv6-delegated-prefix-length** *bits*
no ipv6-delegated-prefix-length

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command allows configuration of delegated prefix length via local user database.

The **no** form of this command reverts to the default.

Parameters *bits* — Specifies the delegated prefix length in bits.

Values 48 to 64

ipv6-delegated-prefix-pool

Syntax **ipv6-delegated-prefix-pool** *pool-name*
no ipv6-delegated-prefix-pool

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command configures the pool name that will be used in DHCPv6 server for DHCPv6 IA-PD prefix selection.

The **no** form of this command removes the pool name from the configuration.

Parameters *pool-name* — Specifies the pool name to be assigned to the delegated prefix pool.

ipv6-lease-times

Syntax [**no**] **ipv6-lease-times**

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command enables the context to configure lease times for DHCPv6.

preferred-lifetime

Syntax **preferred-lifetime** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
preferred-lifetime infinite
no preferred-lifetime

Context config>subscr-mgmt>loc-user-db>ipoe>host>ipv6-lease-times

Description This command specifies the preferred lifetime for an IPv6 lease. When the preferred lifetime expires, then any derived addresses are deprecated.

The **no** form of this command reverts to the default.

Parameters **infinite** — Specifies that the valid lifetime is infinite.

preferred-lifetime — Specifies the preferred lifetime.

Values

days *days*	0 to 3650
hrs *hours*	0 to 23
min *minutes*	0 to 59
sec *seconds*	0 to 59

rebind-timer

Syntax **rebind-timer** [*rebind-timer*]
no rebind-timer

Context config>subscr-mgmt>loc-user-db>ipoe>host>ipv6-lease-times
config>subscr-mgmt>loc-user-db>ppp>host>ipv6-lease-times
config>service>ies>sub-if>grp-if>ipv6>dhcp6>proxy
config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy
config>router>dhcp6>server>pool>prefix
config>service>vprn>dhcp6>server>pool>prefix

Description This command configures the lease rebind timer (T2) via LUDB.

The T2 time is the time at which the client contacts any available addressing authority to extend the lifetimes of DHCPv6 leases. T2 is a time duration relative to the current time expressed in units of seconds.

The IP addressing authority controls the time at which the client contacts the addressing authority to extend the lifetimes on assigned addresses/prefixes through the T1 and T2 parameters assigned to an IA. At time T1 for an IA, the client initiates a Renew/Reply message exchange to extend the lifetimes on any addresses in the IA. The client includes an IA option with all addresses/prefixes currently assigned to the IA in its Renew message. Recommended values for T1 and T2 are .5 and .8 times the shortest preferred lifetime of the addresses/prefixes in the IA that the addressing authority is willing to extend, respectively.

The configured rebind timer should always be longer than or equal to the renew timer.

The T1 and T2 are carried in the IPv6 address option that is within the IA.

The **no** form of this command reverts to the default.

Default rebind-timer min 48

Parameters *rebind-timer* — Specifies the preferred lifetime.

Values

days *days*	0 to 14
hrs *hours*	0 to 23
min *minutes*	0 to 59
sec *seconds*	0 to 9

renew-timer

Syntax **renew-timer** [*renew-timer*]
no renew-timer

Context config>subscr-mgmt>loc-user-db>ipoe>host>ipv6-lease-times
config>subscr-mgmt>loc-user-db>ppp>host>ipv6-lease-times
config>service>ies>sub-if>grp-if>ipv6>dhcp6>proxy
config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy
config>router>dhcp6>server>pool>prefix
config>service>vprn>dhcp6>server>pool>prefix

Description This command configures the lease renew time (T1) via LUDB.

The T1 is the time at which the client contacts the addressing authority to extend the lifetimes of the DHCPv6 leases (addresses/prefixes). T1 is a time duration relative to the current time expressed in units of seconds.

The IP addressing authority controls the time at which the client contacts the addressing authority to extend the lifetimes on assigned addresses through the T1 and T2 parameters assigned to an IA. At time T1 for an IA, the client initiates a Renew/Reply message exchange to extend the lifetimes on any addresses in the IA. The client includes an IA option with all addresses currently assigned to the IA in its Renew message. Recommended values for T1 and T2 are .5 and .8 times the shortest preferred lifetime of the addresses in the IA that the addressing authority is willing to extend, respectively.

The configured renew timer should always be smaller than or equal to the rebind timer.

The T1 and T2 are carried in the IPv6 address option that is within the IA.

The **no** form of this command reverts to the default.

Default renew-timer min 30

Parameters *renew-timer* — Specifies the preferred lifetime.

Values

days *days*	0 to 7
hrs *hours*	0 to 23
min *minutes*	0 to 59
sec *seconds*	0 to 59

valid-lifetime

Syntax **valid-lifetime** [*valid-lifetime*]
valid-lifetime infinite
no valid-lifetime

Context config>subscr-mgmt>loc-user-db>ipoe>host>ipv6-lease-times
config>subscr-mgmt>loc-user-db>ppp>host>ipv6-lease-times
config>service>ies>sub-if>grp-if>ipv6>dhcp6>proxy
config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy

Description This command configured valid-lifetime for DHCPv6 lease (address/prefix).

Valid lifetime is the length of time an address/prefix remains in the valid state (for example., the time until invalidation). The valid lifetime must be greater than or equal to the preferred lifetime. When the valid lifetime expires, the address/prefix becomes invalid and must not be used in communications. RFC 2461, sec 6.2.1 recommends default value of 30 days.

Each address/prefix assigned to the client has associated preferred and valid lifetimes specified by the address assignment authority (DHCP Server, Radius, ESM). To request an extension of the lifetimes assigned to an address, the client sends a Renew message to the addressing authority. The addressing authority sends a Reply message to the client with the new lifetimes, allowing the client to continue to use the address/prefix without interruption.

The lifetimes are transmitted from the addressing authority to the client in the IA option on the top level (not the address or prefix level).

The **no** form of this command reverts to the default.

Default valid-lifetime days 1

Parameters **infinite** — Specifies that the valid lifetime is infinite.

Values 0xffffffff

valid-lifetime — Specifies the preferred lifetime.

Values

days *days*	0 to 3650
hrs *hours*	0 to 23
min *minutes*	0 to 59
sec *seconds*	0 to 9

ipv6-slaac-prefix

Syntax **ipv6-slaac-prefix** *ipv6-prefix/prefix-length*
no ipv6-slaac-prefix

Context config>subscr-mgmt>loc-user-db>ppp>host
config>subscr-mgmt>loc-user-db>ipoe>host

Description This command configures static IPv6 SLAAC prefix (PIO) for the host. The host will assign an IPv6 address to itself based on this prefix. The prefix length is 64 bits.

The **no** form of this command removes the static IPv6 SLAAC prefix (PIO) for the host from the configuration.

Default no ipv6-slaac-prefix

Parameters *ipv6-prefix/prefix-length* — Specifies the IPv6 address and prefix length.

Values

<ipv6-prefix/prefi*> : ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x [0 to FFFF]H
	d [0 to 255]D
prefix-length	64

ipv6-slaac-prefix-pool

Syntax **ipv6-slaac-prefix-pool** *pool*
no ipv6-slaac-prefix-pool

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command configures the IPv6 slaac prefix pool of this host.

The **no** form of this command reverts to the default.

ipv6-wan-address-pool

Syntax **ipv6-wan-address-pool** *pool-name*
no ipv6-wan-address-pool

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command configures the pool name that will be used in DHCPv6 server for DHCPv6 IA-PA address selection.

The **no** form of this command removes the pool name from the configuration.

Parameters *pool-name* — Specifies the WAN address pool, up to 32 characters.

link-address

Syntax **link-address** *ipv6-address*
no link-address

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>service>ies>if>ipv6>dhcp6-relay
config>service>ies>sub-if>ipv6>dhcp6>relay
config>service>ies>sub-if>grp-if>ipv6>dhcp6>relay
config>service>vprn>if>ipv6>dhcp6-relay
config>service>vprn>sub-if>ipv6>dhcp6>relay
config>service>vprn>sub-if>grp-if>ipv6>dhcp6>relay

Description This command configures the link address used for prefix selection at the DHCP server.

The link-address is a field in DHCP6 Relay-Forward message that is used in DHCP6 server to select the IPv6 address (IA-NA) or IPv6 prefix (IA-PD) from a pool with configured prefix range covering the link-address. The selection scope is the pool or a prefix range within the pool.

The **no** form of this command reverts to the default.

Parameters *ipv6-address* — Specifies the link-address.

Values

<ipv6-address>	ipv6-address	- x:x:x:x:x:x:x:x	(eight 16-bit pieces)

x:x:x:x:x:x:d.d.d.d
x - [0 to FFFF]H
d - [0 to 255]D

match-radius-proxy-cache

Syntax **match-radius-proxy-cache**

Context config>subscr-mgmt>loc-user-db>ipoe>host

Description This command enables the context to configure RADIUS proxy cache match parameters.

fail-action

Syntax **fail-action** {**continue** | **drop**}
no fail-action

Context config>subscr-mgmt>loc-user-db>ipoe>host>match-radprox-cache

Description This command specifies the action to take when no match is found in the cache.

The **no** form of this command reverts to the default.

Default fail-action drop

Parameters **continue** — Specifies to continue when no match is found.

drop — Specifies to drop when no match is found.

mac-format

Syntax **mac-format** *mac-format*
no mac-format

Context config>subscr-mgmt>loc-user-db>ipoe>host>match-radprox-cache

Description This command specifies how a MAC address is represented.

Default mac-format “aa:”

Parameters *mac-format* — Specifies the expected MAC address format.

match

Syntax **match** {**circuit-id** | **mac** | **remote-id**}

match option [*number*] [**option6** [*number*]
match option6 [*number*]
no match

Context config>subscr-mgmt>loc-user-db>ipoe>host>match-radprox-cache

Description This command specifies in what DHCPv6 option to retrieve the value to be used as lookup key in the RADIUS proxy cache.

The **no** form of this command reverts to the default.

Default match mac

Parameters **circuit-id** — Specifies to use the circuit Id to match against.

mac — Specifies the MAC address to match against.

remote-id — Specifies the remote ID to match against.

option *number* — Specifies the option number that the DHCP server uses to send the identification strings to the client.

Values 1 to 124

option6 *number* — Specifies the DHCPv6 option to retrieve the value to be used as lookup key in the RADIUS Proxy cache.

Values 1 to 65535

server

Syntax **server** [**service** *service-id*] **name** *server-name*]
no server

Context config>subscr-mgmt>loc-user-db>ipoe>host>match-radprox-cache

Description This command specifies the RADIUS proxy server.

The **no** form of this command reverts to the default.

Parameters *service-id* — Specifies an existing service ID or service name.

Values *service-id* - 1 to 214748364
service-name - up to 64 characters.

server-name — Specifies the name of local user database server, up to 32 characters.

mld-parameters

Syntax **mld-parameters**

Context config>subscr-mgmt>loc-user-db>ppp>host

config>subscr-mgmt>loc-user-db>ipoe>host

Description This command enables the context to configure an MLD import policy.

import

Syntax **import** *policy-name*
no import

Context config>subscr-mgmt>loc-user-db>ppp>host>mld-parameters
config>subscr-mgmt>loc-user-db>ipoe>host>mld-parameters

Description This command configures an MLD import policy.

The LUDB allows a list of up to 14 MLD import policies per host. The MLD policy also allows the configuration of an additional import policy, providing a total of 15 MLD import policies per host. The import policy inside the MLD policy is always applied last, which determines if the list is a black list or a white list. To configure an MLD white list, the import policies in the LUDB should all be allow or forward entries and the import policy in the MLD policy should have a default action to deny all. To configure a black list, the import policies inside the LUDB should be drop entries and the MLD policy import policy default action should be to forward all. The 15 import policies can be configured to be a mixed white and black list. Since it is difficult to control the order of the import policies within the LUDB, the recommendation is to provision the import policy inside the MLD policy first for deterministic behavior.

The **no** version of this command removes the specified import policy.

Parameters *policy-name* — Specifies the MLD import policy used to control the multicast group accessible for the subscriber host.

msap-defaults

Syntax **msap-default**

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command configures MSAP authentication defaults.

group-interface

Syntax **group-interface** *ip-int-name* [**prefix** {*port-id*}] [**suffix** {*port-id*}]
no group-interface

Context config>subscr-mgmt>loc-user-db>ipoe>host>msap-defaults
config>subscr-mgmt>loc-user-db>ppp>host>msap-defaults

Description This command configures the group interface.

The **no** form of this command removes the group interface parameters from the configuration.

Parameters *ip-int-name* — Specifies the IP interface name, up to 32 characters.

port-id — Specifies the port ID as the prefix to the specified IP interface name.

port-id — Specifies the port ID as the suffix to the specified IP interface name.

policy

Syntax **policy** *msap-policy-name*
no policy

Context config>subscr-mgmt
config>subscr-mgmt>loc-user-db>ipoe>host>msap-defaults
config>subscr-mgmt>loc-user-db>ppp>host>msap-defaults

Description This command configures the MSAP policy.

The **no** form of this command removes the MSAP policy name from the configuration.

Parameters *msap-policy-name* — Specifies the policy name, up to 32 characters.

service

Syntax **service** *service-id*
no service

Context config>subscr-mgmt>loc-user-db>ipoe>host>msap-defaults
config>subscr-mgmt>loc-user-db>ppp>host>msap-defaults

Description This command sets default service for all subscribers created based on trigger packets received on the given capture SAP in case the corresponding VSA is not included in the RADIUS authentication response. This command is applicable to capture SAP only.

The **no** form of this command reverts to the default.

Parameters *service-id* — Specifies the service ID as an integer or the service name.

Values *service-id* - 1 to 2147483648
service-name - up to 64 characters

custom-option

Syntax **custom-option** *option-number* **address** [*ip-address*]

custom-option *option-number* **hex** *hex-string*
custom-option *option-number* **string** *ascii-string*
no custom-option *option-number*

Context config>router>dhcp>server>pool>options
config>router>dhcp>server>pool>subnet>options
config>subscr-mgmt>loc-user-db>ipoe>host>options
config>subscr-mgmt>loc-user-db>ppp>host>options
config>router>dhcp6>server>pool>options
config>router>dhcp6>server>pool>prefix>options
config>service>vprn>dhcp>server>pool>options
config>service>vprn>dhcp6>server>pool>options
config>service>vprn>dhcp6>server>pool>prefix>options

Description This command configures specific DHCP options. The options defined here can overrule options in the local user database.

The **no** form of the removes the custom option parameters from the configuration.

Parameters *option-number* — Specifies up to four option numbers that the DHCP server uses to send the identification strings to the DHCP client.

Values 1 to 254

ip-address — Specifies the IP address of this host.

hex-string — Specifies the hex value of this option.

Values 0x0 to 0xFFFFFFFF (maximum 254 hex nibbles)

ascii-string — Specifies the value of this option, up to 127 characters.

default-router

Syntax **default-router** *ip-address* [*ip-address*]
no default-router

Context config>router>dhcp>server>pool>subnet>options
config>subscr-mgmt>loc-user-db>ipoe>host>options
config>service>vprn>dhcp6>server>pool>subnet

Description This command configures the IP address of the default router for a DHCP client. Up to four IP addresses can be specified.

The **no** form of this command removes the address(es) from the configuration.

Parameters *ip-address* — Specifies up to four default router IP addresses. Each address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

dns-server

Syntax **dns-server** *ip-address* [*address*]
no dns-server

Context config>router>dhcp>server>pool>options
config>subscr-mgmt>loc-user-db>ipoe>host>options
config>subscr-mgmt>loc-user-db>ppp>host>options
config>service>vprn>dhcp>server>pool>options

Description This command configures the IPv6 address of the DNS server.

The **no** form of this command removes the IPv6 address from the DNS server configuration.

Parameters *ipv6-address* — Specifies up to four DNS server IPv6 addresses. Each address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

domain-name

Syntax **domain-name** *domain-name*
no domain-name

Context config>router>dhcp>server>pool>options
config>subscr-mgmt>loc-user-db>ipoe>host>options
config>router>dhcp6>server>pool>options
config>router>dhcp6>server>pool>prefix>options
config>service>vprn>dhcp>server>pool>options
config>service>vprn>dhcp6>server>pool>options
config>service>vprn>dhcp6>server>pool>prefix>options

Description This command configures the default domain for a DHCP client that the router uses to complete unqualified host names (without a dotted-decimal domain name).

The **no** form of this command removes the name from the configuration.

Parameters *domain-name* — Specifies the domain name for the client, up to 127 characters.

lease-rebind-time

Syntax **lease-rebind-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no lease-rebind-time

Context config>router>dhcp>server>pool>options
config>subscr-mgmt>loc-user-db>ipoe>host>options
config>service>vprn>dhcp>server>pool>options

Description This command configures the time the client transitions to a rebinding state.

The **no** form of this command removes the time from the configuration.

Parameters *lease-rebind-time* — Specifies the lease rebind time.

Values

days *days*	0 to 3650
hrs *hours*	0 to 23
min *minutes*	0 to 59
sec *seconds*	0 to 59

lease-renew-time

Syntax **lease-renew-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no lease-renew-time

Context config>router>dhcp>server>pool>options
config>subscr-mgmt>loc-user-db>ipoe>host>options
config>service>vprn>dhcp>server>pool>options

Description This command configures the time the client transitions to a renew state.

The **no** form of this command removes the time from the configuration.

Parameters *lease-renew-time* — Specifies the lease renew time.

Values

days:	0 to 3650
hours:	0 to 23
minutes:	0 to 59
seconds	0 to 59

lease-time

Syntax **lease-time** [*lease-time*]
no lease-time

Context config>router>dhcp>server>pool>options
config>subscr-mgmt>loc-user-db>ipoe>host>options
config>service>vprn>dhcp>server>pool>options

Description This command configures the amount of time that the DHCP server grants to the DHCP client permission to use a particular IP address.

The **no** form of this command removes the lease time parameters from the configuration.

Parameters *lease-time* — Specifies the lease time.

Values

days *days*	0 to 3650
hrs *hours*	0 to 23
min *minutes*	0 to 59
sec *seconds*	0 to 59

netbios-name-server

Syntax **netbios-name-server** *ip-address* [*ip-address*]
no netbios-name-server

Context config>router>dhcp>server>pool>options
config>subscr-mgmt>loc-user-db>ppp>host>options
config>subscr-mgmt>loc-user-db>ipoe>host>options
config>service>vprn>dhcp>server>pool>options

Description This command configures up to four Network Basic Input/Output System (NetBIOS) name server IP addresses.

The **no** form of this command removes the IP address from the netbios-name-server configuration.

Parameters *ip-address* — Specifies up to four NetBIOS name server IP addresses. The address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

netbios-node-type

Syntax **netbios-node-type** *netbios-node-type*
no netbios-node-type

Context config>router>dhcp>server>pool>options
config>subscr-mgmt>loc-user-db>ipoe>host>options
config>service>vprn>dhcp>server>pool>options

Description This command configures the Network Basic Input/Output System (NetBIOS) node type.

The **no** form of this command removes the netbios node type parameters from the configuration.

Parameters *netbios-node-type* — Specifies the netbios node type.

Values B — Broadcast node uses broadcasting to query nodes on the network for the owner of a NetBIOS name.

P — Peer-to-peer node uses directed calls to communicate with a known NetBIOS name server for the IP address of a NetBIOS machine name.

M — Mixed node uses broadcast queries to find a node, and if that fails, queries a known P-node name server for the address.

H — Hybrid node is the opposite of the M-node action so that a directed query is executed first, and if that fails, a broadcast is attempted.

subnet-mask

Syntax **subnet-mask** *ip-address*
no subnet-mask

Context config>router>dhcp>server>pool>subnet>options
config>subscr-mgmt>loc-user-db>ipoe>host>options

Description This command specifies the subnet-mask option to the client. The mask can either be defined (for supernetting) or taken from the pool address.

The **no** form of this command removes the address from the configuration.

Parameters *ip-address* — Specifies the IP address of the subnet mask. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

options6

Syntax **options6**

Context config>subscr-mgmt>loc-user-db>ppp>host
config>subscr-mgmt>loc-user-db>ipoe>host

Description This command enables the context to configure IPv6 DNS server information in the local user database

dns-server

Syntax **dns-server** *ipv6-address* [*ipv6-address*]
no dns-server

Context config>router>dhcp6>server>pool>options
config>router>dhcp6>server>pool>prefix>options
config>subscr-mgmt>loc-user-db>ppp>host>options6
config>subscr-mgmt>loc-user-db>ipoe>host>options6

config>service>vprn>dhcp6>server>pool>prefix>options

Description This command configures IPv6 DNS server addresses that can be used for name resolution

The **no** form of this command removes the IPv6 address from the dns-server configuration.

Parameters *ipv6-address* — Specifies IPv6 DNS server addresses.

retail-service-id

Syntax **retail-service-id** *service-id*
no retail-service-id

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command indicates the service ID of the retailer VPRN service to which this session belongs. If the value of this object is non-zero, the session belongs to a retailer VPRN.

The **no** form of this command removes the service ID from the configuration.

Parameters *service-id* — Specifies the retailer service ID or name.

Values

service-id:	1 to 2147483647
service-name:	Service name up to 64 characters.

router-advertisement-policy

Syntax **router-advertisement-policy** *policy*
no router-advertisement-policy

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command applies an RA policy to the host.

The **no** form of this command removes the policy from the configuration.

Parameters *policy* — Specifies the policy name, up to 32 characters.

server6

Syntax **server6** *ipv6-address*
no server6

Context config>subscr-mgmt>loc-user-db>ipoe>host

Description This command allows DHCP6 server selection based on the host entry in LUDB.

The configured DHCP6 server IP address must reference one of the v6 addressees configured under the **config>service>vprn>sub-if>grp-if>ipv6>dhcpv6>relay** or **config>service>ies>sub-if>grp-if>ipv6>dhcpv6>relay** context.

The **no** form of this command reverts to the default.

Parameters *ipv6-address* — Specifies the retailer service ID.

Values

ipv6-address	ipv6-address - x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D

to-client-options

Syntax **to-client-options**

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command enables the context to configure DHCP options to send to the client.

option

Syntax **option** *option-number* **address** *ipv4-address* [*ipv4-address*]
option *option-number* **hex** *hex-string*
option *option-number* **string** *ascii-string*
no option *option-number*

Context config>subscr-mgmt>loc-user-db>ipoe>host>to-client-options>ipv4
config>subscr-mgmt>loc-user-db>ppp>host>to-client-options>ipv4

Description This command configures DHCPv4 options via LUDB that are passed in all DHCP messages to the client. The options are blindly appended to any options already present in the DHCP message. In other words, there is no intelligent merging of the options where overlapping options from different sources are further evaluated to determine whether only one option or multiple options should be returned to the client.

Multiple DHCP options can be configured at the same time although each option requires its own option statement. Those options are equivalent to RADIUS VSAs **Alc-ToCLient-Dhcp4-Options**.

DHCPv4 options can be provided via DHCPv4 server in the relay case. In addition, DHCPv4 options provided via LUDB or RADIUS can be supplied and consequently appended to the already existing options. If DHCPv4 options are provided simultaneously via LUDB and RADIUS, the RADIUS as a source of DHCPv4 option is blocked and the options supplied via LUDB are passed to the client. This is valid for the relay and proxy case.

Any DHCP option may be encoded in the option statement. An example of the option statement for DHCPv4 default-gateway is given below:

option 3 192.168.1.254

DHCPv4 options may be fixed length or variable length. They are appended at the end of DHCPv4 messages. All options begin with a tag octet, which uniquely identifies the option. Fixed-length options without data consist of only a tag octet. Only options 0 and 255 are fixed length. All other options are variable-length.

The **no** form of the removes the option from the configuration.

Parameters *option-number* — Specifies up to four option numbers. This can be a well-known or a an anonymous option.

ipv4-address — Specifies an IPv4 address as an option.

hex-string — Specifies options coded as Hex characters.

ascii-string — Specifies options coded as an ASCII string.

to-server-options

Syntax [**no**] **to-server-options**

Context config>subscr-mgmt>loc-user-db>ipoe>host

Description This command enables the context to configure DHCP options to send to the server

option

Syntax **option** *option-number* **address** [*ip-address*]
no option *option-number*

Context config>subscr-mgmt>loc-user-db>ipoe>host>to-server-options>ipv6

Description This command specifies the DHCPv6 options to send to the server. There are four options:

- **option** *option-number* **hex** *hex-string*
- **option** *option-number* **string** *ascii-string*
- **option** *option-number* **domain** d*omain-name*
- **no option** *option-number*

The **no** form of this command removes the option parameters from the configuration.

Parameters *option-number* — Specifies the option number and the identification string that will inserted in the DHCP client message to the server.

Values 1 to 65535

ip-address — Specifies the IPv6 address.

hex-string — Specifies the hex value of the option.

Values 0x0 to 0xFFFFFFFF (maximum 254 hex nibbles)

ascii-string — Specifies the string value of the option.

Values 127 character max

domain-name — Specifies the domain as in RFC 1035 section 3.1.

wpp

Syntax **wpp**

Context config>subscr-mgmt>loc-user-db>ipoe>host

Description This command enables the context to configure WPP parameters.

initial-app-profile

Syntax **initial-app-profile** *app-profile-name*
no initial-app-profile

Context config>subscr-mgmt>loc-user-db>ipoe>host>wpp
config>service>ies>sub-if>grp-if>wpp
config>service>vprn>sub-if>grp-if>wpp

Description This command specifies the initial app-profile for the hosts created on the group-interface. This initial app-profile is replaced after hosts pass web portal authentication.

The **no** form of this command reverts to the default.

Parameters *app-profile-name* — Specifies the initial application profile, up to 32 characters, to be used during the WPP authentication phase of the IPoE hosts.

initial-sla-profile

Syntax **initial-sla-profile** *sla-profile-name*
no initial-sla-profile

Context config>subscr-mgmt>loc-user-db>ipoe>host>wpp

config>service>ies>sub-if>grp-if>wpp
config>service>vprn>sub-if>grp-if>wpp

Description This command specifies the initial sla-profile for the hosts created on the group-interface. This initial sla-profile is replaced after hosts pass web portal authentication.

The **no** form of this command reverts to the default.

Parameters *sla-profile-name* — Specifies the initial SLA profile, up to 32 characters, to be used during the WPP authentication phase of the IPOE host.

initial-sub-profile

Syntax **initial-sub-profile** *sub-profile-name*
no initial-sub-profile

Context config>subscr-mgmt>loc-user-db>ipoe>host>wpp
config>service>ies>sub-if>grp-if>wpp
config>service>vprn>sub-if>grp-if>wpp

Description This command specifies the initial sub-profile for the hosts created on the group-interface. This initial sub-profile is replaced after hosts pass web portal authentication.

The **no** form of this command reverts to the default.

Parameters *sub-profile-name* — Specifies the initial subscriber profile, up to 32 characters, to be used during the WPP authentication phase of the IPoE host.

portal

Syntax [**no**] **portal router** *router-instance* **name** *wpp-portal-name*

Context config>subscr-mgmt>loc-user-db>ipoe>host>wpp

Description This command configures the portal for the portal group. This command is mutually exclusive with the **portal-group** command.

Parameters *router-instance* — Specifies the virtual router instance.

Values

router-name:	Base, management
service-id:	1 to 2147483647
service-name:	Specifies the service name, up to 64 characters

Default Base

wpp-portal-name — Specifies the name of the web portal server, up to 32 characters.

portal-group

Syntax **portal-group** *portal-group-name*
no portal-group

Context config>subscr-mgmt>loc-user-db>ipoe>host>wpp

Description This command configures the WPP portal group name. This command is mutually exclusive with the **portal** command.

Parameters *portal-group-name* — Specifies the WPP portal group name, up to 32 characters.

restore-disconnected

Syntax [**no**] **restore-disconnected**

Context config>service>ies>sub-if>grp-if>wpp
config>service>vprn>sub-if>grp-if>wpp

Description This command specifies that the initial profiles must be restored after a DHCP host has disconnected. The behavior that system will restore the **initial-sla-profile**, **initial-sub-profile**, or **initial-app-profile** when hosts disconnects instead of removing them.

The **no** form of this command specifies that the initial profiles will not be restored after a DHCP host has disconnected.

Default restore-disconnected

mask

Syntax **mask type** *match-type* {[**prefix-string** *prefix-string* | **prefix-length** *prefix-length*] [**suffix-string** *suffix-string* | **suffix-length** *suffix-length*]}
no mask type *match-type*
mask type *dhcp-match-type* {[**prefix-string** *prefix-string* | **prefix-length** *prefix-length*] [**suffix-string** *suffix-string* | **suffix-length** *suffix-length*]}
no mask type *dhcp-match-type*

Context config>subscr-mgmt>loc-user-db>ppp
config>subscr-mgmt>loc-user-db>ipoe

Description This command configures the mask.

The **no** form of this command removes the mask parameters from the configuration.

Parameters *ppp-match-type* — Specifies the sub-option inserted by the PPPoE intermediate agent.

Values circuit-id, remote-id, sap-id, service-name, username

dhcp-match-type — The data type represents the type of matching done to identify a DHCP host.

Values circuit-id, option60, remote-id, sap-id, string, system-id

prefix-string — Specifies a substring that is stripped of the start of the incoming circuit ID before it is matched against the value configured in the DHCP or PPPoE circuit ID.

This string can only contain printable ASCII characters. The “*” character is a wildcard that matches any substring. If a "\" character is masked, use the escape key so it becomes "\\".

Values 127 characters maximum, “*”

prefix-length — Specifies the number of characters to remove from the start of the incoming circuit ID before it is matched against the value configured in the circuit ID

Values 1 to 127

suffix-string — Specifies a substring that is stripped of the end of the incoming circuit ID before it is matched against the value configured in circuit ID.

This string can only contain printable ASCII characters. The “*” character is a wildcard that matches any substring. If a "\" character is masked, use the escape key so it becomes "\\".

Values 127 characters maximum

suffix-length — Specifies the number of characters to remove from the end of the incoming circuit ID before it is matched against the value configured in the circuit ID.

Values 1 to 127

match-list

Syntax **match-list** *match-type-1* [*match-type-2*]
no match-list

Context config>subscr-mgmt>loc-user-db>ppp
config>subscr-mgmt>loc-user-db>ipoe

Description This command specifies the type of matching done to identify a host. There are different match-types for PPPoE or IPoE hosts of which a maximum of four can be specified.

The **no** form of this command reverts to the default.

Parameters *match-type-x* — Specifies up to four matching types to identify a host.

Values circuit-id, derived-id, dual-stack-remote-id, encap-tag-range, ip, mac, option60, remote-id, sap-id, service-id, string, system-id

circuit-id — Specifies the DHCP4 option (82,1) or DHCP6 option 18.

mac — Specifies the MAC address of the client. Chaddr in DHCP4 and DUID in IPv6.

option60 — Specifies the DHCP4 option 60.

remote-id — Specifies the DHCP4 option (82,2) or DHCP6 option 37 .

Note: The format of remote-id in IPv6 is different that the format of remote-id in IPv4; IPv6 remote-id contains enterprise-id filed that is also honored in matching.

dual-stack-remote-id — Specifies the enterprise-id in v6 Remote-id will be stripped off before LUDB matching is performed. Processing of IPv4 Remote-id remains unchanged. This will allow a single host entry in LUDB for dual-stack host where host identification is performed based on the Remote-id field.

sap-id — Specifies the SAP ID on which DHCPv4 packet are received. The sap-id is inserted as ALU VSO (82,9,4) by the DHCPv4 relay in router. This is enabled via configuration under the vendor-specific-option CLI hierarchy of the DHCPv4 relay. Since the dhcp-relay configuration is enabled under the group-interface CLI hierarchy, the group-interface and the service-id must be known before the sap-id can be used for LUDB match.

encap-tag-range — Specifies the VLAN tags.

ip — Specifies the source IPv4/IPv6 address of a data-trigger packet.

service-id — Specifies the service-id of the ingress SAP for DHCPv4 packets. The service-id is inserted as ALU VSO (82,9,3) by the DHCPv4 relay in router. This is enabled via configuration under the vendor-specific-option CLI hierarchy of the DHCPv4 relay.

string — Specifies the custom string configured under the vendor-specific-option CLI hierarchy of the DHCPv4 relay. The string is inserted as ALU VSO (82,9,5) by the DHCPv4 relay in router. Since the dhcp-relay configuration is enabled under the group-interface CLI hierarchy, the group-interface and the service-id must be known before the string can be used for LUDB match.

system-id — Specifies the system-id of the node name configured under the **system>name** CLI hierarchy. The system-id is inserted as ALU VSO (82,9,1) by the DHCPv4 relay in router. This is enabled via configuration under the vendor-specific-option CLI hierarchy of the DHCPv4 relay. Since the dhcp-relay configuration is enabled under the group-interface CLI hierarchy, the group-interface and the service-id must be known before the system-id can be used for LUDB match.

derived-id — Specifies the value extracted by Python script during processing of DHCP Discover/Solicit/Request/Renew/Rebind Messages (client to server bound messages). The value is stored in the DHCP Transaction Cache (DTC) in a variable named alc.dtc.derivedId. This value has a lifespan of a DHCP transaction (a single pair of messages exchanged between the client and the server; for example DHCP Discover and DHCP Offer).

ppp

Syntax **ppp**

Context config>subscr-mgmt>loc-user-db

Description This command enables the context to configure PPP host parameters.

access-loop-encapsulation

Syntax [**no**] **access-loop-encapsulation**

Context config>subscr-mgmt>loc-user-db>ppp>host

Description This command enables the context to configure access loop information.

encap-offset

Syntax **encap-offset** [**type** *type*]
no encap-offset

Context config>subscr-mgmt>loc-user-db>ppp>host>ali

Description This command configures the egress encapsulation offset.

The **no** form of this command reverts to the default.

Parameters **type** *type* — Selects the encap type.

Values pppoa-llc, pppoa-null, pppoeoa-llc, pppoeoa-llc-fcs, pppoeoa-llc-tagged, pppoeoa-llc-tagged-fcs, pppoeoa-null, pppoeoa-null-fcs, pppoeoa-null-tagged, pppoeoa-null-tagged-fcs, ipoa-llc, ipoa-null, ipoeoa-llc, ipoeoa-llc-fcs, ipoeoa-llc-tagged, ipoeoa-llc-tagged-fcs, ipoeoa-null, ipoeoa-null-fcs, ipoeoa-null-tagged, ipoeoa-null-tagged-fcs, pppoe, pppoe-tagged, ipoe, ipoe-tagged

rate-down

Syntax **rate-down** *rate*
no rate-down

Context config>subscr-mgmt>loc-user-db>ppp>host>ale

Description This command configures the last mile link downstream rate in the access loop.

The **no** form of this command reverts to the default.

Parameters *rate* — Specifies the last mile link downstream rate needed for proper (shaping) rate calculations and interleaving delay in the access loop.

Values 1 to 100000 kb/s

access-loop-information

Syntax **access-loop-information**

Context config>subscr-mgmt>loc-user-db>ppp>host

Description This command enables the context to configure access loop information in the local user database

circuit-id

Syntax **circuit-id sap-id**
circuit-id string *ASCII string*
no circuit-id

Context config>subscr-mgmt>loc-user-db>ppp>host>ali

Description This command specifies a circuit-id for PPPoE hosts. A circuit ID received in PPPoE tags has precedence over the LUDB specified circuit ID.

The **no** form of this command reverts to the default.

Parameters **sap-id** — Specifies to use the SAP ID of the PPPoE session as the circuit ID.

ASCII string — Specifies the circuit ID as a string, up to 63 characters.

remote-id

Syntax **remote-id mac**
remote-id string *ASCII string*
no remote-id

Context config>subscr-mgmt>loc-user-db>ppp>host>ali

Description This command specifies a remote-id for PPPoE hosts. A remote-id received in PPPoE tags has precedence over the LUDB specified remote ID.

The **no** form of this command reverts to the default.

Parameters **mac** — Specifies MAC address of the PPPoE session as the remote ID.

ASCII string — Specifies the circuit ID as a string, up to 63 characters.

force-ipv6cp

Syntax [**no**] **force-ipv6cp**

Context config>subscr-mgmt>loc-user-db>ppp>host

Description This command specifies if the IPv6 control protocol should be negotiated after PPP reaches the Network-Layer Protocol phase.

The **no** form of this command reverts to the default.

l2tp

Syntax **l2tp**

Context config>subscr-mgmt>loc-user-db>ppp>host

Description This command enables the context to configure L2TP parameters for the host.

group

Syntax **group** *tunnel-group-name* [**service-id** *service-id*]
no group

Context config>subscr-mgmt>loc-user-db>ppp>host>l2tp

Description This command configures the L2TP tunnel group. The tunnel-group-name is configured in the **config>router>l2tp** context. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Router Configuration Guide*.

Parameters *tunnel-group-name* — Specifies an existing tunnel L2TP group, up to 63 characters.

service-id *service-id* — Specifies an existing service ID or service name.

Values *service-id*: 1 to 214748364
service-name: A string, up to 64 characters

authentication-policy

Syntax **authentication-policy** *policy-name*
no authentication-policy

Context config>subscr-mgmt>loc-user-db>ppp>host

Description This command configures the authentication policy for the host. A host name with name “default” will be matched when all other hosts do not match.

padi-auth-policy

Syntax **padi-auth-policy** *policy-name*

no padi-auth-policy

Context config>subscr-mgmt>loc-user-db>ppp>host

Description This command configures the PADI authentication policy of this host.

Parameters *policy-name* — Specifies the authentication policy name.

pado-delay

Syntax **pado-delay** *deci-seconds*
no pado-delay

Context config>subscr-mgmt>loc-user-db>ppp>host

Description This command configures the delay timeout before sending a PPPoE Active Discovery Offer (PADO).

Parameters *deci-seconds* — Specifies the delay timeout before sending a PADO.

Values 1 to 30

password

Syntax **password** {**ignore** | **chap** *string* |**pap** *string*} [**hash**|**hash2**]
no password

Context config>subscr-mgmt>loc-user-db>ppp>host

Description This command specifies a password type or configures password string for **pap** or **chap**. The pap and chap passwords are stored in a hashed format in the config files. The **hash|hash2** optional keywords are used for config execution.

This command will only be interpreted if the local user database is connected directly to the PPPoE node under the VPRN/IES group interface. It is not used if the local user database is accessed by a local DHCP server.

The **no** form of this command reverts to the default.

Parameters **ignore** — Specifies that the password will be ignored, in which case authentication will always succeed, independent of the password used by the PPPoE client. The client must still perform authentication.

string — Specifies that the password for Challenge-Handshake Authentication Protocol) (CHAP) is used. Only a password received with the CHAP protocol will be accepted.

string — Specifies that the Password Authentication Protocol (PAP) is used. Only a password received with the PAP protocol will be accepted, even though the CHAP protocol will be proposed to the client first because it is unknown at the time of the offer which password type will be allowed to the client.

hash|hash2 — Specifies hashing scheme

pre-auth-policy

Syntax **pre-auth-policy** *policy-name*
no pre-auth-policy

Context config>subscr-mgmt>loc-user-db>ppp>host

Description This command configures the RADIUS pre-authentication policy to use to authenticate the PPP host.

The **no** form of this command reverts to the default.

Parameters *policy-name* — Specifies the pre-authentication policy of the host, up to 32 characters.

steering-profile

Syntax **steering-profile** *steering-profile-name*
no steering-profile

Context config>subscr-mgmt>loc-user-db>ppp>host

Description This command configures the steering profile for the specific host.

The **no** form of this command removes the steering profile for the host.

Parameters *steering-profile-name* — Specifies the name of the steering profile, up to 32 characters.

ignore-df-bit

Syntax [**no**] **ignore-df-bit**

Context config>subscr-mgmt>loc-user-db>ppp>host

Description When the ignore-df-bit command is enabled for a subscriber host, then the do-not-fragment (DF) bit in the IPv4 header for frames egressing the subscriber interface is ignored: the frames are fragmented according the applicable egress MTU; the DF bit is reset for frames that are fragmented.

This command applies to PPPoE PTA and L2TP LNS frames only. It is not applicable for L2TP LAC frames.

interface

Syntax **interface** *ip-int-name* **service-id** *service-id*
no interface

Context config>subscr-mgmt>loc-user-db>ppp>host

Description This command configures the interface where PPP sessions are terminated.

The **no** form of this command reverts to the default.

Parameters *ip-int-name* — Specifies the name of the group interface where the PPP sessions are established.

service-id — Specifies the service ID or name of the service where the PPP sessions are established.

Values *service-id* - 1 to 2147483647
service-name - up to 64 characters

server-id

Syntax **server-id duid-en hex** *hex-string*
server-id duid-en string *ascii-string*
server-id duid-ll
no server-id

Context config>service>ies>sub-if>grp-if>ipv6>dhcp6>proxy
config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy
config>router>dhcp6>server
config>service>vprn>dhcp6>server

Description This command allows operator to customize the “server-id” attribute of a DHCPv6 message (such as DHCPv6 advertise and reply). By default, the server-id uses DUID-ll derive from the chassis link-layer address. Operators have the option to use a unique identifier by using DUID-en (vendor based on enterprise number). There is a maximum length associated with the customizable hex-string and ascii-string.

The **no** form of this command reverts to the default.

Default server-id duid-ll

Parameters *hex-string* — Specifies a DUID system ID in a hex format.

Values 0x0 to 0xFFFFFFFF (max 116 hex nibbles

ascii-string — Specifies a DUID system ID in an ASCII format, up to 58 characters.

duid-ll — Specifies that the DUID system ID is derived from the system link layer address.

on-link

Syntax [**no**] **on-link**

Context config>router>router-advert>if
config>service>ies>sub-if>ipv6>rtr-adv>pfx-opt
config>service>ies>sub-if>grp-if>ipv6>rtr-adv>pfx-opt
config>service>vprn>router-advert>if
config>service>vprn>sub-if>ipv6>rtr-adv>pfx-opt
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv>pfx-opt

Description This command specifies whether the prefix is assigned to an interface on the specified link.

The **no** form of this command reverts to the default.

server-id

Syntax **server-id duid-en hex** *hex-string*
server-id duid-en string *ascii-string*
server-id duid-ll
no server-id

Context config>router>dhcp6>server

Description This command allows operator to customize the “server-id” attribute of a DHCPv6 message from the DHCPv6 proxy server (such as DHCPv6 advertise and reply). By default, the server-id uses DUID-ll derive from the chassis link-layer address. Operators have the option to use a unique identifier by using DUID-en (vendor based on enterprise number). There is a maximum length associated with the customizable hex-string and ascii-string.

The **no** form of this command reverts to the default.

Default server-id duid-ll

Parameters *hex-string* — Specifies a DUID system ID in a hex format.

Values 0x0 to 0xFFFFFFFF (max 116 hex nibbles

ascii-string — Specifies a DUID system ID in an ASCII format, up to 58 characters.

duid-en — Specifies the enterprise-number.

duid-ll — Specifies that the DUID system ID is derived from the system link layer address.

3.6.2.1.5 Router DHCP Commands

local-dhcp-server

Syntax **local-dhcp-server** *server-name* [**create**]
no local-dhcp-server *server-name*

Context config>router>dhcp
config>service>vprn>dhcp

Description This command instantiates a local DHCP server. A local DHCP server can serve multiple interfaces but is limited to the routing context it was which it was created.

The **no** form of this command reverts to the default.

Parameters *server-name* — Specifies the name of local DHCP server

create — Keyword used to create the local DHCP server. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

allow-lease-query

Syntax [**no**] **allow-lease-query**

Context config>router>dhcp6>server
config>service>vprn>dhcp6>server

Description If enabled, the local DHCPv6 server will handle and reply to lease query messages.

The **no** form of this command disables lease query support.

failover

Syntax **failover**

Context config>router>dhcp>server
config>service>vprn>dhcp>server

Description This command enables the context to configure failover parameters.

ignore-mclt-on-takeover

Syntax [**no**] **ignore-mclt-on-takeover**

Context config>router>dhcp>server>failover
config>router>dhcp>server>pool>failover

config>router>dhcp6>server>failover
config>router>dhcp6>server>pool>failover
config>service>vprn>dhcp>server>failover
config>service>vprn>dhcp>server>pool>failover
config>service>vprn>dhcp6>server>failover
config>service>vprn>dhcp6>server>pool>failover

Description With this flag enabled, the remote IP address/prefix can be taken over immediately upon entering the PARTNER-DOWN state of the intercommunication link, without having to wait for the Maximum Client Lead Time (MCLT) to expire. By setting this flag, the lease times of the existing DHCP clients, while the intercommunication link is in the PARTNER-DOWN state, will still be reduced to the MCLT over time and all new lease times are set to MCLT. This behavior remains the same as originally intended for MCLT.

Some deployments require that the remote IP address/prefix range starts delegating new IP addresses/prefixes upon the failure of the intercommunication link, without waiting for the intercommunication link to transition from the COMM-INT state into the PARTNER-DOWN state and the MCLT to expire while in PARTNER-DOWN state.

This can be achieved by enabling the **ignore-mclt-on-takeover** flag and by configuring the **partner-down-delay** to 0.

Enabling this functionality must be exercised with caution. One needs to keep in mind that the partner-down-delay and MCLT timers were originally introduced to prevent IP address duplication in cases where DHCP redundant nodes transition out-of-sync due to the failure of intercommunication link. These timers (**partner-down-delay** and MCLT) would ensure that during their duration, the new IP addresses/prefixes are delegated only from one node, the one with local IP address-range/prefix. The drawback is of course that the new IP address delegation is delayed and thus service is impacted.

But if one could ensure that the intercommunication link is always available, then the DHCP nodes would stay in sync and the two timers would not be needed. This is why it is of utmost importance that in this mode of operation, the intercommunication link is well protected by providing multiple paths between the two DHCP nodes. The only event that should cause intercommunication link to fail is the entire nodal failure. This failure is acceptable since in this case only one DHCP node is available to provide new IP addresses/prefixes.

The **no** form of this command reverts to the default.

maximum-client-lead-time

Syntax **maximum-client-lead-time** [*maximum-client-lead-time*]
no maximum-client-lead-time

Context config>router>dhcp>server>failover
config>router>dhcp>server>pool>failover
config>service>vprn>dhcp>server>failover
config>service>vprn>dhcp>server>pool>failover
config>router>dhcp6>server>failover

config>router>dhcp6>server>pool>failover
config>service>vprn>dhcp6>server>failover
config>service>vprn>dhcp6>server>pool>failover

Description The **maximum-client-lead-time** (MCLT) is the maximum time that a DHCP server can extend client's lease time beyond the lease time currently known by the DHCP partner node. In dual-homed environment, the initial lease time for all DHCP clients is by default restricted to MCLT. Consecutive DHCP renews are allowed to extend the lease time beyond the MCLT.

The MCLT is a safeguard against IP address/prefix duplication in cases of a lease synchronization failure when local-remote failover model is deployed

Once the intercommunication link failure between the redundant DHCP servers is detected, the DHCP IP address range configured as remote will not be allowed to start delegating new leases until the MCLT + partner-down-delay intervals expire. This is to ensure that the new lease that was delegated from the 'local' IP address-range/prefix on one node, but was never synchronized due to the intercommunication link failure, will expire before the same IP address/prefix is allocated from the remote IP address-range/prefix on the other node.

However, the already existing (and synchronized) lease times can be renewed from the remote IP address range at any time, regardless of the state of the intercommunication link (operational or failed).

Lease synchronization failure can be caused either by a node failure, or a failure of the link over which the DHCP leases are synchronized (intercommunication link). Synchronization failure detection can take up to 3 seconds.

During the failure, the DHCP lease time for the new clients is restricted to MCLT while for the existing clients the lease time will over time (by consecutive DHCP renews) be gradually reduced to the MCLT.

The **no** form of this command reverts to the default.

Default maximum-client-lead-time min 10

Parameters *maximum-client-lead-time* — Specifies the maximum client lead time

Values

hrs *hours*	1 to 23
min *minutes*	1 to 59
sec *seconds*	1 to 59

partner-down-delay

Syntax **partner-down-delay** [*partner-down-delay*]
no partner-down-delay

Context config>router>dhcp>server>failover
config>router>dhcp>server>pool>failover

config>service>vprn>dhcp>server>failover
config>service>vprn>dhcp>server>pool>failover
config>router>dhcp6>server>failover
config>router>dhcp6>server>pool>failover
config>service>vprn>dhcp6>server>failover
config>service>vprn>dhcp6>server>pool>failover

Description Since the DHCP lease synchronization failure can be caused by the failure of the intercommunication link (and not necessary the entire node), there is a possibility the redundant DHCP servers become isolated in the network. In other words, they can serve DHCP clients but they cannot synchronize the lease. This can lead to duplicate assignment of IP addresses, since the servers have configured overlapping IP address ranges but they are not aware of each other's leases.

The purpose of the partner-down-delay is to prevent the IP lease duplication during the intercommunication link failure by not allowing new IP addresses to be assigned from the remote IP address range. This timer is intended to provide the operator with enough time to remedy the failed situation and to avoid duplication of IP addresses/prefixes during the failure.

During the partner-down-delay time, the prefix designated as remote is eligible only for renewals of the existing DHCP leases that have been synchronized by the peering node. Only after the sum of the partner-down-delay and the maximum-client-lead-time will the prefix designated as remote be eligible for delegation of the new DHCP leases. When this occurs, we say that the remote IP address range has been taken over.

It is possible to expedite the takeover of a remote IP address range so that the new IP leases can start being delegated from that range shortly after the intercommunication failure is detected. This can be achieved by configuring the partner-down-delay timer to 0 seconds, along with enabling the ignore-mclt-on-takeover CLI flag. Caution must be taken before enabling this functionality. It is safe to bypass safety timers (partner-down-delay + MCLT) only in cases where the operator is certain that the intercommunication between the nodes has failed due to the entire node failure and not due to the intercommunication (MCS) link failure. Failed intercommunication due to the nodal failure would ensure that only one node is present in the network for IP address delegation (as opposed to two isolated nodes with overlapping IP address ranges where address duplication can occur). For this reason, the operator must ensure that there are redundant paths between the nodes to ensure uninterrupted synchronization of DHCP leases.

In access-driven mode of operation, partner-down-delay has no effect.

The **no** form of this command reverts to the default.

Default partner-down-delay hrs 23 min 59 sec 59

Parameters *partner-down-delay* — Specifies the partner down delay time

Values

hrs *hours*	1 to 23
min *minutes*	1 to 59
sec *seconds*	0 to 59

peer

Syntax **peer** *ip-address* **tag** *sync-tag*
no peer

Context config>router>dhcp>server>failover
config>router>dhcp>server>pool>failover
config>router>dhcp6>server>failover
config>router>dhcp6>server>pool>failover

Description DHCP leases can be synchronized per DHCP server of DHCP pool. The pair of synchronizing servers or pools is identified by a tag. The synchronization information is carried over the Multi-Chassis Synchronization (MCS) link between the two peers. MCS link is a logical link (IP, or MPLS).

MCS runs over TCP, port 45067 and it is using either data traffic or keepalives to detect failure on the communication link between the two nodes. In the absence of any MCS data traffic for more than 0.5sec, MCS will send its own keepalive to the peer. If a reply is **not** received within three sec, MCS will declare its operation state as DOWN and the DB Sync state as out-of-sync. MCS will consequently notify its clients (DHCP Server being one of them) of this. It can take up to three seconds before the DHCP client realizes that the inter-chassis communication link has failed.

The inter-chassis communication link failure does not necessarily assume the same failed fate for the access links. The two redundant nodes can become isolated from each other in the network. This would occur in cases where only the intercommunication (MCS) link fails. It is of utmost importance that this MCS link be highly redundant.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the IPv4 address of the peer

tag *sync-tag* — Specifies a tag that will identify synchronizing DHCP servers or pools

startup-wait-time

Syntax [**no**] **startup-wait-time** [**min** *minutes*] [**sec** *seconds*] [**hrs** *hours*]

Context config>router>dhcp>server>failover
config>router>dhcp>server>pool>failover
config>router>dhcp6>server>failover
config>router>dhcp6>server>pool>failover

Description This command enables startup-wait-time during which each peer waits after the initialization process before assuming the active role for the prefix designated as local or access-driven. This is to avoid transient issues during the initialization process.

The **startup-wait-time** should be configured to an interval in which, after boot, both nodes can set up an MCS TCP link and start MCS. The timer is restarted each time the server downloads a lease from the MCS database and stops when the last state record from the peer is synchronized. The next state is (PRE-)NORMAL, unless the timer times out or is forced to stop via a new tools command (tools perform router dhcp or **dhcp6 local-dhcp-server** *server-name* [**pool**] **failover abort-startup-wait**), in which case the local DHCP server will transition immediately to the COMMUNICATIONS-INTERRUPTED state.

Default startup-wait-time min 2

Parameters *minutes* — Specifies the startup wait time, in minutes

Values 1 to 59

seconds — Specifies the startup wait time, in seconds

Values 1 to 59

hours — Specifies the start wait time, in hours

Values 1

force-renews

Syntax [**no**] **force-renews**

Context config>router>dhcp>server
config>service>vprn>server

Description This command enables the sending of sending FORCERENEW messages.

The **no** form of this command disables the sending of FORCERENEW messages.

ignore-rapid-commit

Syntax [**no**] **ignore-rapid-commit**

Context config>router>dhcp6>server
config>service>vprn>dhcp6>server

Description This command enables the Rapid Commit Option.

The **no** form of this command disables the Rapid Commit Option.

interface-id-mapping

Syntax [**no**] **interface-id-mapping**

Context config>router>dhcp6>server

config>service>vprn>dhcp6>server

Description If enabled, this command enables the behavior where unique /64 prefix is allocated per interface-id, and all clients having the same interface-id get an address allocated out of this /64 prefix. This is relevant for bridged clients behind the same local-loop (and same SAP), where sharing the same prefix allows communication between bridged clients behind the same local-loop to stay local. For SLAAC based assignment, downstream neighbor-discovery is automatically enabled to resolve the assigned address.

The **no** form of this command reverts to the default.

lease-hold-time

Syntax **lease-hold-time** [*lease-hold-time*]
no lease-hold-time

Context config>service>vprn>dhcp>server
config>router>dhcp>server
config>service>vprn>dhcp6>server
config>router>dhcp6>server

Description This command configures the time to remember this lease and is applicable for unsolicited release conditions such as lease timeout if the **lease-hold-time-for** command is set to the default value **no solicited-release** and is additionally applicable for normal solicited releases from DHCP clients if the **lease-hold-time-for** command is set to **solicited-release**.

The **no** form of this command reverts to the default.

Default lease-hold-time sec 0

Parameters *lease-hold-time* — Specifies the amount of time to remember the lease

Values

days *days*	0 to 7305
hrs *hours*	0 to 23
min *minutes*	0 to 59
sec *seconds*	0 to 59

lease-hold-time-for

Syntax [**no**] **lease-hold-time-for**

Context config>service>vprn>dhcp6>server
config>router>dhcp6>server
config>service>vprn>dhcp>server
config>router>dhcp>server

Description This command enables the context to configure **lease-hold-time-for** parameters which define additional types of lease or triggers that cause system to hold up leases.

The **no** form of this command reverts to the default.

internal-lease-ipsec

Syntax [**no**] **internal-lease-ipsec**

Context config>service>vprn>dhcp6>server>lease-hold-time-for
config>router>dhcp6>server>lease-hold-time-for
config>service>vprn>dhcp>server
config>router>dhcp>server

Description This command enables the server to hold up the lease of local IPsec clients.

The **no** form of this command disables the ability of the server to hold up the lease of local IPsec clients.

solicited-release

Syntax [**no**] **solicited-release**

Context config>service>vprn>dhcp6>server>lease-hold-time-for
config>router>dhcp6>server>lease-hold-time-for
config>service>vprn>dhcp>server
config>router>dhcp>server

Description This command enables the server to hold up a lease even in case of solicited release; for example, when the server receives a normal DHCP release message.

The **no** form of this command disables the ability of the server to hold up a lease when a solicited release is received.

pool

Syntax **pool** *pool-name* [**create**]
no pool *pool-name*

Context config>router>dhcp>server
config>router>dhcp6>server
config>service>vprn>dhcp>server
config>service>vprn>dhcp6>server

Description This command configures a DHCP address pool on the router.

The **no** form of this command removes the pool name from the configuration.

Parameters *pool name* — Specifies the name of this IP address pool. Allowed values are any string, up to 32 characters.

create — Keyword used to create the pool. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

delegated-prefix-length

Syntax **delegated-prefix-length** *bits* [**minimum** *prefix-length*] [**maximum** *prefix-length*]
no delegated-prefix-length

Context config>router>dhcp6>server>pool
config>service>vprn>dhcp6>server>pool

Description This command configures the delegated prefix length that is used if the DHCPv6 client does not specify a prefix length hint.

The DHCPv6 client prefix length hint is limited by the range specified by the **minimum** and **maximum** parameters. If the hint is smaller than the minimum, the allocated prefix length is equal to the minimum length. If the hint is larger than the maximum, the allocated prefix length is equal to the maximum length.

The **no** form of this command reverts to the default.

Default delegated-prefix-length 64 minimum 48 maximum 64

Parameters *bits* — Specifies the delegated prefix length, in bits

Values 48 to 127

prefix-length — Specifies the minimum or maximum allowed prefix length, in bits

Values 48 to 127

exclude-prefix

Syntax [**no**] **exclude-prefix** *ipv6-prefix/prefix-length*

Context config>router>dhcp6>server>pool
config>service>vprn>dhcp6>server>pool

Description This command defines a prefix that to be excluded from available prefix in the pool. The typical use case is to exclude the interface address.

- A held lease is deleted if it got excluded by an exclude prefix.
- An exclude range can never exclude only a part of an existing lease. If for example a /63 PD is assigned, an exclude of /64 which belongs to this /63 can not be configured.
- A single exclude prefix can never exclude a whole include prefix.

- When applying or removing an exclude prefix, the threshold stats are adjusted to reflect the actual address space and its usage.

The **no** form of this command removes the prefix that is to be excluded from available prefix in the pool.

Parameters *ipv6-prefix/prefix-length* — Specifies an IPv6 prefix and prefix length

Values

ipv6-prefix	x:x:x:x:x:x:x:x	(eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x - [0 to FFFF]H
		d - [0 to 255]D
	prefix-length	[0 to 128]

max-lease-time

Syntax **max-lease-time** [*max-lease-time*]
no max-lease-time

Context config>router>dhcp>server>pool

Description This command configures the maximum lease time.

The **no** form of this command reverts to the default.

Default max-lease-time days 10

Parameters *time* — Specifies the maximum lease time

Values

days *days*	0 to 3650
hrs *hours*	0 to 23
min *minutes*	0 to 59
sec *seconds*	0 to 59

min-lease-time

Syntax **min-lease-time** [*min-lease-time*]
no min-lease-time

Context config>router>dhcp>server>pool
config>service>vprn>dhcp>server>pool

Description This command configures the minimum lease time.

The **no** form of this command reverts to the default.

Default min-lease-time min 10

Parameters *time* — Specifies the minimum lease time

Values

days *days*	0 to 3650
hrs *hours*	0 to 23
min *minutes*	0 to 59
sec *seconds*	0 to 59

minimum-free

Syntax **minimum-free** *minimum-free* [**percent**] [**event-when-depleted**]
no minimum-free

Context config>router>dhcp>server>pool
config>router>dhcp6>server>pool

Description This command specifies the desired minimum number of free addresses in this pool.

The **no** form of this command reverts to the default.

Default minimum-free 1

Parameters *minimum-free* — Specifies the minimum number of free addresses

Values 0 to 255

percent — Specifies that the value indicates a percentage

event-when-depleted — This parameter enables a system-generate event when all available addresses in the pool/subnet of local DHCP server are depleted.

nak-non-matching-subnet

Syntax [**no**] **nak-non-matching-subnet**

Context config>router>dhcp>server>pool
config>service>vprn>dhcp>server>pool

Description With this command, if the local DHCPv4 server receives a DHCP request with option 50 (means client try to request a previous allocated message as described in section 3.2 of RFC 2131, *Dynamic Host Configuration Protocol*) and the address allocation algorithm ends up using a pool and the address in option50 is not in pool, then system will return a DHCP NAK, otherwise system just drop the DHCP packet.

The **no** form of this command reverts to the default.

Default no nak-non-matching-subnet

offer-time

Syntax **offer-time** [**min** *minutes*] [**sec** *seconds*]
no offer-time

Context config>router>dhcp>server>pool
config>service>vprn>dhcp>server>pool

Description This command configures the time interval during which a DHCP offer is valid.

The **no** form of this command reverts to the default.

Default offer-time min 1

Parameters *time* — Specifies the offer time

Values

min *minutes*	0 to 10
sec *seconds*	0 to 59

options

Syntax **options**

Context config>router>dhcp>server>pool
config>router>dhcp>server>pool>subnet
config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host
config>router>dhcp6>server>pool
config>router>dhcp6>server>pool>prefix
config>service>vprn>dhcp>server>pool>
config>service>vprn>dhcp6>server>pool>
config>service>vprn>dhcp6>server>pool>prefix>

Description This command enables the context to configure pool options. The options defined here can be overruled by defining the same option in the local user database.

retail-service-id

Syntax [**no**] **retail-service-id** *service-id*

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command sets the retail-service ID for a given subscriber host.

The **no** form of this command reverts to the default.

Parameters *service-id* — Specifies the service ID as an integer or service name, up to 64 characters

Values 1 to 2147483647, svc-name

prefix

Syntax **prefix** *ipv6-addr/prefix-length* [**failover** {**local** | **remote** | **access-driven**}] [**pd**] [**wan-host**] [**create**]
no prefix *ipv6-addr/prefix-length*

Context config>router>dhcp6>server>pool
config>service>vprn>dhcp6>server>pool

Description This command allocates a prefix to a pool from which Prefix Delegation prefixes and or wan addresses can be assigned.

The **no** form of this command removes the prefix parameters from the configuration.

Parameters *prefix ipv6-addr/prefix-length* — Specifies the prefix

Values

ipv6-address	x:x:x:x:x:x:x:x	(eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x [0 to FFFF]H
		d [0 to 255]D
	prefix-length	[1 to 128]

failover {**local** | **remote | access-driven**} — This command designates a prefix as local, remote, or access-driven. This is used when multi-chassis synchronization is enabled.

Values **local** — (default) A prefix designated as local is always used to renew the existing addresses/prefixes or to assign a new one.

remote — A prefix designated as remote is used only to renew the existing DHCP leases. The new leases are assigned from it only after the **maximum-client-lead-time** and **partner-down-delay** *time* elapses.

access-driven — A prefix designated as access-driven is similar to local (a new prefix assignment as well as a renewal). However, as the prefix is shared between the redundant server pair, the following additional conditions should be met to avoid duplicate address allocations:

- A dual home access protection mechanism such as srrp or mc-lag must ensure a single active path from the dhcp client to the server.
- The dhcp relay should point to the local server only.

pd — (default) When configured, IA-PD prefixes may be allocated from this prefix.

wan-host — (default) When configured, IA-NA addresses and/or SLAAC prefixes (local-address-assignment) may be allocated from this prefix

create — Keyword used to create the prefix configuration. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

thresholds

Syntax **thresholds**

Context config>service>vprn>dhcp6>server>pool
config>service>vprn>dhcp6>server>pool>prefix
config>router>dhcp6>server>pool
config>router>dhcp6>server>pool>prefix

Description This command enables the context to configure pool level thresholds.

Default thresholds

minimum-free

Syntax [**no**] **minimum-free prefix-length** [*prefix-length*]

Context config>service>vprn>dhcp6>server>pool>thresholds
config>router>dhcp6>server>pool>thresholds

Description This command creates a threshold for a given prefix length on the pool level. Up to 128 thresholds could be created. For example, with **minimum-free prefix-length 64**, then the usage of /64 prefix in the pool is counted.

There are two types of thresholds that could be defined at the pool level:

- Depleted — The system sends out a warning when the prefix with the configured length is no long available in the pool.
- Minimum free — A percentage-based threshold which represents the minimal available percentage of prefix with the configured length in the pool. The system will send out warning if the actual percentage is lower than the configured percentage

Configuration of this command also enables the system stats collection for **configure prefix length**, which could be displayed via the **show router** *router-id* **dhcp6 local-dhcp-server d6 pool-threshold-stats** command.

The **no** form of this command removes the prefix-length from the configuration.

Parameters *prefix-length* — Specifies the IPv6 prefix length

Values 1 to 128

depleted-event

Syntax [**no**] **depleted-event**

Context config>service>vprn>dhcp6>server>pool>thresholds>minimum-free
config>router>dhcp6>server>pool>thresholds>minimum-free

Description This command enables the system to send out warnings when the prefix with the configured length is no long available in the pool.

The **no** form of this command disables the warnings.

depleted-event

Syntax [**no**] **depleted-event**

Context config>service>vprn>dhcp6>server>pool>prefix>thresholds>minimum-free
config>router>dhcp6>server>pool>prefix>thresholds>minimum-free

Description This command enables the system to send out a warning when the prefix with a configured length is no long available in the provisioned prefix.

For example:

```
prefix 2001:0:0:ffe0::/50 pd wan-host create
   thresholds
       minimum-free prefix-length 64
           depleted-event
```

With the above configuration, the system will send out a warning when there is no available /64 that can be allocated out of 2001:0:0:ffe0::/50.

The **no** form of this command disables the warnings.

minimum

Syntax **minimum percent** [*percent*] [**number** [*number*]]
no minimum

Context config>service>vprn>dhcp6>server>pool>thresholds>minimum-free
config>router>dhcp6>server>pool>thresholds>minimum-free

Description This command specifies a percentage based threshold which represent the minimal available percentage of the prefix with configured length in the pool. The system will send out a warning if the actual percentage is lower than the configured percentage.

The **no** form of this command removes the percentage from the configuration.

Parameters *percent* — Specifies the percentage of used prefixes with the minimum free threshold length in the pool compared to the number of provisioned prefixes

Values 0 to 100

number — Specifies the number of used prefixes with the minimum free threshold length in the pool compared to the number of provisioned prefixes

Values 0 to 4294967295

minimum

Syntax **minimum** [**percent** [*percent*]] [**number** [*number*]]
no minimum

Context config>service>vprn>dhcp6>server>pool>prefix>thresholds>minimum-free
config>router>dhcp6>server> pool>prefix>thresholds>minimum-free

Description This command configures a percentage-based or number-based threshold which represents the minimal available percentage or number of the prefix with a configured length in the provisioned prefix. The system will send out a warning if the actual percentage or number is lower than the configured threshold.

For example:

```
prefix 2001:0:0:ffe0::/50 pd wan-host create
   thresholds
       minimum-free prefix-length 64
           minimum number 3
```

With the above configuration, the system will send a warning when the number of available /64 in prefix 2001:0:0:ffe0::/50 is less than 3.

The **no** form of this command removes the command parameters from the configuration.

Parameters *percent* — Specifies the percentage of used prefixes with the minimum free threshold length in the pool compared to the number of provisioned prefixes

Values 0 to 100

number — Specifies the number of prefixes

Values 0 to 4294967295

user-db

Syntax **user-db** *name*
no user-db

Context config>subscr-mgmt>loc-user-db>ppp>host

Description This command specifies local user database for PPP PAP/CHAP access.

With this configuration, system will access the specified DB again during PPP PAP/CHAP phase.

This configuration only becomes effective when system is accessing parent DB during PPPoE discovery phase.

The **no** form of this command removes the name from the configuration.

Parameters *name* — Specifies the name of the local user database for PAP/CHAP access

custom-option

Syntax **custom-option** *option-number* **address** [*ip-address*]
custom-option *option-number* **domain** [*domain-string*]
custom-option *option-number* **hex** *hex-string*
custom-option *option-number* **string** *ascii-string*
no custom-option *option-number*

Context config>router>dhcp>server>pool>options
config>router>dhcp6>server>pool>options
config>router>dhcp>server>pool>subnet>options
config>service>vprn>dhcp>server>pool>options
config>service>vprn>dhcp6>server>pool>options

Description This command configures specific DHCP options. The options defined here can overrule options in the local user database.

The **no** form of the removes the option from the configuration.

Parameters *option-number* — Specifies up to four DHCP server option numbers used to send the identification strings to the DHCP client

Values 1 to 254

ip-address — Specifies the IP address of this host

domain-string — Specifies the domain string, up to 127 characters for DHCP6 only.

hex-string — Specifies the hex value of this option

Values 0x0 to 0xFFFFFFFF (maximum 254 hex nibbles)

ascii-string — Specifies the value of this option, up to 127 characters

subnet

Syntax **subnet** {*ip-address*/*mask* | *ip-address netmask*} [**create**]
no subnet {*ip-address*/*mask* | *ip-address netmask*}

Context config>router>dhcp>server>pool
config>service>vprn>dhcp>server>pool

Description This command creates a subnet of IP addresses to be served from the pool. The subnet cannot include any addresses that were assigned to subscribers without those addresses specifically excluded. When the subnet is created no IP addresses are made available until a range is defined.

The **no** form of the removes the subnet parameters from the configuration.

Parameters *ip-address* — Specifies the base IP address of the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

mask — The subnet mask in dotted decimal notation. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252.

Note: A mask of 255.255.255.255 is reserved for system IP addresses.

netmask — Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains

create — Keyword used to create the subnet. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

address-range

Syntax [**no**] **address-range** *start-ip-address end-ip-address* [**failover** {**local** | **remote**}]

Context config>router>dhcp>server>pool>subnet
config>service>vprn>dhcp>server>pool

Description This command configures a range of IP addresses to be served from the pool. All IP addresses between the start and end IP addresses will be included (other than specific excluded addresses).

The **no** form of this command removes the address-range parameters from the configuration.

Parameters *start-ip-address* — Specifies the start address of this range to include. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

end-ip-address — Specifies the end address of this range to include. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

local — Specifies that the DHCP server failover control type is in control under normal operation

remote — Specifies that the remote DHCP server failover system is in control under normal operation

drain

Syntax [**no**] **drain**

Context config>service>vprn>dhcp>server>pool>subnet

Description This command means no new leases can be assigned from this subnet and existing leases are cleaned up upon renew/rebind.

The **no** form of this command means the subnet is active and new leases can be assigned from it.

exclude-addresses

Syntax [**no**] **exclude-addresses** *start-ip-address* [*end-ip-address*]

Context config>router>dhcp>server>pool>subnet
config>service>vprn>dhcp>server>pool>subnet

Description This command specifies a range of IP addresses that excluded from the pool of IP addresses in this subnet.

The **no** form of the removes the parameters from the configuration.

Parameters *start-ip-address* — Specifies the start address of this range to exclude. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

end-ip-address — Specifies the end address of this range to exclude. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

maximum-declined

Syntax **maximum-declined** *maximum-declined*
no maximum-declined

Context config>router>dhcp>server>pool>subnet
config>service>vprn>dhcp>server>pool>subnet

Description This command configures the maximum number of declined addresses allowed.

The **no** form of the reverts to the default.

Default maximum-declined 64

Parameters *maximum-declined* — Specifies the maximum number of declined addresses allowed

Values 0 to 4294967295

minimum-free

Syntax **minimum-free** *minimum-free* [**percent**] [**event-when-depleted**]
no minimum-free

Context config>router>dhcp>server>pool>subnet
config>router>dhcp6>server>pool>subnet
config>service>vprn>dhcp>server>pool>subnet
config>service>vprn>dhcp6>server>pool>subnet

Description This command configures the minimum number of free addresses in this subnet. If the actual number of free addresses in this subnet falls below this configured minimum, a notification is generated.

The **no** form of the reverts to the default.

Default minimum-free 1

Parameters *minimum-free* — Specifies the minimum number of free addresses in this subnet

Values 0 to 255

percent — Specifies that the value indicates a percentage

event-when-depleted — Enables a system-generate event when all available addresses in the pool/subnet of local DHCP server are depleted.

use-gi-address

Syntax **use-gi-address** [**scope** *scope*]

Context config>router>dhcp>server
config>service>vprn>dhcp>server

Description This command enables the use of gi-address matching. If the gi-address flag is enabled, a pool can be used even if a subnets is not found. If the **local-user-db-name** is not used, the gi-address flag is used and addresses are handed out by GI only. If a user must be blocked from getting an address the server maps to a local user database and configures the user with no address.

A pool can include multiple subnets. Since the GI is shared by multiple subnets in a subscriber interface the pool may provide IP addresses from any of the subnets included when the GI is matched to any of its subnets. This allows a pool to be created that represents a sub-int.

The **no** form of the reverts to the default.

Parameters *scope* — Specifies if addresses are handed out for a certain subnet where the gi-address belongs to only or for all subnets part of the pool

Values **subnet** — Addresses are only handed out for the subnet where the gi-address is part.

pool — All subnets part of the pool which contain subnet where the gi-address is part of can hand out addresses.

use-pool-from-client

Syntax **use-pool-from-client delimiter** *delimiter*
use-pool-from-client
no use-pool-from-client

Context config>router>dhcp>server
config>service>vprn>dhcp>server

Description This command enables the use of the pool indicated by DHCP client. When enabled, the IP address pool to be used by this server is the pool is indicated by the vendor-specific sub-option 13 of the DHCP option 82. When disabled or if there is no sub-option 13 in the DHCP message, the pool selection falls back to the **use-gi-address** configuration.

The **no** form of the command disables the use of the pool indicated by DHCP client.

Parameters *delimiter* — A single ASCII character specifies the delimiter of separating primary and secondary pool names in Option82 VSO

user-ident

Syntax **user-ident** *user-ident*
no user-ident

Context config>router>dhcp>server
config>service>vprn>dhcp>server

Description This command configures the user identification method for the DHCPv4 server.

The **no** form of the reverts to the default.

Default user-ident mac-circuit-id

Parameters *user-ident* — Specifies the user identification method

Values
- **client-id** — Specifies to use the DHCPv4 client identifier as the user identification method.
- **circuit-id** — Specifies to use the circuit identifier of the DHCPv4 client as the user identification method.
- **mac** — Specifies to use the MAC address of the DHCPv4 client as the user identification method.
- **mac-circuit-id** — Specifies to use the MAC address and circuit identifier of the DHCPv4 client as the user identification method.
- **remote-id** — Specifies to use the MAC address of the remote end as the user identification method.

user-ident

Syntax **user-ident** *user-ident*
no user-ident

Context config>router>dhcp6>server
config>service>vprn>dhcp6>server

Description This command configures the keys for identification of the DHCPv6 lease being held in the lease-database (for configured period after lease timeout). Subscriber requesting a lease via DHCPv6 that matches an existing lease based on this configured key is handed the matched prefix or address. This allows address and prefix “stickiness” for DHCPv6 assigned prefixes (IA_NA or PD).

The **no** form of the reverts to the default.

Default user-ident duid

Parameters *user-ident* — Specifies the user identification method

Values
- **duid** — Specifies the IPv6 DHCP unique identifier from DHCPv6.
- **interface-id** — Specifies the IPv6 interface-id option.
- **interface-id-link-local** — Specifies the interface-id and link-local address.

use-link-address

Syntax **use-link-address** [**scope** *scope*]
no use-link-address

Context config>router>dhcp6>server
config>service>vprn>dhcp6>server

Description This command configures the local pool selection for v6 address or prefix assignment to use the configured link-address under relay configuration. The selected pool will contain a prefix covering the link-address. The scope option defines the scope for the match. With scope **subnet**, the prefix or address selection is limited to the prefix in the pool that covers the link-address. With scope **pool**, all the prefixes in the selected pool are eligible for assignment.

The **no** form of the reverts to the default.

Default scope subnet

Parameters *scope* — Specifies the scope of the IP address selection

Values **subnet** — Specifies that the prefix or address selection is limited to the prefix in the pool that covers the link address.

pool — Specifies that all prefixes in the selected pool are eligible for assignment.

user-db

Syntax **user-db** *local-user-db-name* [**create**]
no user-db

Context config>router>dhcp>server
config>subscr-mgmt
config>service>ies>sub-if>grp-if>dhcp
config>service>vprn>sub-if>grp-if>dhcp

Description This command configures a local user database for authentication.

The **no** form of this command reverts to the default.

Parameters *local-user-db-name* — Specifies the name of a user database, up to 32 characters

create — Keyword used to create the user database. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

3.6.2.1.6 DHCP Service Commands

dhcp

Syntax **dhcp**

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp
config>service>ies>interface

config>service>vprn
config>service>vprn>if
config>service>vprn>sub-if
config>service>vprn>sub-if>grp-if
config>service>ies>sub-if>grp-if
config>service>ies>sub-if
config>service>ies>sub-if>grp-if

Description This command enables the context to configure DHCP parameters.

anti-spoof

Syntax **anti-spoof** *type*
no anti-spoof

Context config>service>ies>sap
config>service>vpls>sap
config>service>ies>sub-if>gr-if>pppoe
config>service>ies>ies>sub-if>gr-if>sap
config>service>ies>ies>sub-if>gr-if>sap-parameters
config>service>vprn>sub-if>grp-if>pppoe
config>service>vprn>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>sap-parameters
config>subscr-mgmt>msap-policy>ies-vprn-only-sap-parameters
config>service>vpls>sap

Description This command enables anti-spoof filtering and optionally changes the anti-spoof matching type for the SAP.

The type of anti-spoof filtering defines what information in the incoming packet is used to generate the criteria to lookup an entry in the anti-spoof filter table. The type parameter (**ip**, **mac**, **ip-mac**) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled.

The **no** form of this command disables anti-spoof filtering on the SAP.

Parameters *type* — Specifies the anti-spoof filtering type for this SAP

Values ip — Specifies SAP anti-spoof filtering to use only the source IP address in its lookup. If a static host exists on the SAP without an IP address specified, the **anti-spoof ip** command fails.

mac — Specifies SAP anti-spoof filtering to use only the source MAC address in its lookup. If a static host exists on the SAP without a specified MAC address, the **anti-spoof mac** command fails.

ip-mac — Specifies SAP anti-spoof filtering to use both the source IP address and the source MAC address in its lookup.

circuit-id

Syntax **circuit-id** [**ascii-tuple** | **ifindex** | **sap-id** | **vlan-ascii-tuple**]
no circuit-id

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option

Description When enabled, the router sends an ASCII-encoded tuple in the **circuit-id** sub-option of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by "|".

If disabled, the **circuit-id** sub-option of the DHCP packet is left empty.

The **no** form of this command returns the system to the default.

Default circuit-id ascii-tuple

Parameters **ascii-tuple** — Specifies that the ASCII-encoded concatenated tuple be used which consists of the access-node-identifier, service-id, and interface-name, separated by "|"

ifindex — Specifies that the interface index is used. (The If index of a router interface can be displayed using the **show>router>interface>detail** command)

sap-id — Specifies that the SAP identifier is used

vlan-ascii-tuple — Specifies that the format includes VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q-encapsulated ports only. Thus, when the Option 82 bits are stripped, the dot1p bits are copied to the Ethernet header of an outgoing packet.

dhcp6

Syntax **dhcp6**

Context config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp
config>service>ies>interface
config>system
config>service>vprn
config>service>vprn>interface
config>service>vprn>sub-if
config>service>vprn>sub-if>grp-if
config>service>vprn>sub-if>grp-if>ipv6
config>service>ies>sub-if>grp-if>ipv6
config>service>ies>sub-if
config>service>ies>sub-if>grp-if

Description This command enables the context to configure DHCP6 parameters.

client-applications

Syntax **client-applications** {[**dhcp**] [**ppp**]}
no client-applications

Context config>service>vprn>sub-if>dhcp
config>service>ies>sub-if>dhcp
config>service>vprn>sub-if>grp-if>dhcp
config>service>ies>sub-if>grp-if>dhcp
config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy
config>service>vprn>sub-if>grp-if>ipv6>dhcp6>relay
config>service>ies>sub-if>grp-if>ipv6>dhcp6>proxy
config>service>ies>sub-if>grp-if>ipv6>dhcp6>relay
config>service>vprn>sub-if>ipv6>dhcp6>proxy
config>service>vprn>sub-if>ipv6>dhcp6>relay
config>service>ies>sub-if>ipv6>dhcp6>proxy
config>service>ies>sub-if>ipv6>dhcp6>relay

Description This command enables DHCP relay and proxy-server for the configured client types.

The **no** form of this command reverts to the default.

Default dhcp

Parameters **dhcp** — Enables IPoE clients to use the DHCP relay or proxy-server

ppp — Enables PPPoE clients to use the DHCP relay or proxy-server that PPPoE will attempt to request an IP address for a PPPoE client from the DHCP server(s) assigned to PPPoE node

match-circuit-id

Syntax [**no**] **match-circuit-id**

Context config>service>ies>sub-if>grp-if>dhcp
config>service>vprn>sub-if>dhcp
config>service>vprn>sub-if>grp-if>dhcp

Description This command enables Option 82 circuit ID on relayed DHCP packet matching. For routed CO, the group interface DHCP relay process is stateful. When packets are relayed to the server the virtual router ID, transaction ID, SAP ID, and client hardware MAC address of the relayed packet are tracked.

When a response is received from the server the virtual router ID, transaction ID, and client hardware MAC address must be matched to determine the SAP on which to send the packet out. In some cases, the virtual router ID, transaction ID, and client hardware MAC address are not guaranteed to be unique.

When the **match-circuit-id** command is enabled this as part of the key is used to guarantee correctness in our lookup. This is really only needed when dealing with an IP aware DSLAM that proxies the client hardware MAC address.

The **no** form of this command disables Option 82 circuit ID on relayed DHCP packet matching.

lease-populate

Syntax **lease-populate** [*nbr-of-entries*]
no lease-populate

Context config>service>ies>if>dhcp
config>service>ies>if>dhcp>option
config>service>ies>sub-if>grp-if>dhcp
config>service>vpls>sap>dhcp
config>service>vpls>if>dhcp>option
config>service>vprn>if>dhcp
config>service>vprn>sub-if>dhcp
config>service>vprn>sub-if>grp-if>dhcp
config>subscr-mgmt>msap-policy>vpls-only>dhcp

Description This command enables dynamic host lease state management for SAPs.

For VPLS, DHCP snooping must be explicitly enabled (using the **snoop** command) at all points where DHCP messages requiring snooping enter the VPLS instance (both from the DHCP server and from the subscribers). Lease state information is extracted from snooped DHCP ACK messages to populate lease state table entries for the MSAP.

The optional number-of-entries parameter is used to define the number lease state table entries allowed for an MSAP or IP interface. If number-of-entries is omitted, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed and subsequent DHCP ACK messages are discarded.

The retained lease state information representing dynamic hosts may be used to:

- Populate an MSAP based anti-spoof filter table to provide dynamic anti-spoof filtering. If the system is unable to populate the dynamic host information in the anti-spoof filter table on the SAP, the DHCP ACK message must be discarded without adding new lease state entry or updating an existing lease state entry.
- Generate dynamic ARP replies if **arp-reply-agent** is enabled.

The **no** form of this command disables dynamic host lease state management for the MSAP.

lease-populate

Syntax **lease-populate** [*nbr-of-leases*]

lease-populate [*nbr-of-leases*] **l2-header** [**mac** *ieee-address*]
no lease-populate

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp
config>service>vpls>sap>dhcp
config>service>ies>if>dhcp
config>service>vprn>if>dhcp
config>service>ies>sub-if>grp-if>dhcp
config>service>vprn>sub-if>grp-if>dhcp
config>service>vprn>sub-if>dhcp

Description This command enables and disables dynamic host DHCPv4 lease state management for SAPs.

For VPLS, DHCP snooping must be explicitly enabled (using the **snoop** command) at all points where DHCP messages requiring snooping enter the VPLS instance (both from the DHCP server and from the subscribers). Lease state information is extracted from snooped DHCP ACK messages to populate lease state table entries for the SAP.

The optional number-of-entries parameter defines the number lease state table entries allowed.

- for this SAP in case of a VPLS service
- for this interface in case of an IES or VPRN interface
- for each SAP in case of an IES or VPRN group-interface
- for this interface in case of an IES or VPRN retail subscriber-interface

If number-of-entries is omitted, only a single entry is allowed. Once the maximum number of entries has been reached, subsequent lease state entries are not allowed and subsequent DHCP ACK messages are discarded.

The retained lease state information representing dynamic hosts may be used to:

- Populate a SAP based anti-spoof filter table to provide dynamic anti-spoof filtering. If the system is unable to populate the dynamic host information in the anti-spoof filter table on the SAP, the DHCP ACK message must be discarded without adding new lease state entry or updating an existing lease state entry.
- Populate the system’s ARP cache based on the arp-populate configuration. Applicable to IES and VPRN interfaces or group-interfaces.
- Populate managed entries into a VPLS forwarding database. VPLS forwarding database population is an implicit feature that automatically places the dynamic host’s MAC address into the VPLS FDB. When a dynamic host’s MAC address is placed in the lease state table, it will automatically be populated into the VPLS forwarding database associated with the SAP on which the host is learned. The dynamic host MAC address will override any static MAC entries using the same MAC and prevent dynamic learning

of the MAC on another interface. Existing static MAC entries with the same MAC address as the dynamic host are marked as inactive but not deleted. If all entries in the lease state table associated with the MAC address are removed, the static MAC may be populated. New static MAC definitions for the VPLS instance may be created while a dynamic host exists associated with the static MAC address

- Generate dynamic ARP replies if **arp-reply-agent** is enabled. Applicable to VPLS service SAPs

The **no** form of this command reverts to the default.

Parameters *nbr-of-leases* — Specifies the number of DHCPv4 leases allowed

l2-header — Indicates a mode of operation where anti-spoof entry associated with the given DHCP state is created based on the MAC address from the Layer 2 header. The Layer 2 header flag is not set by default. This parameter is only applicable for group interfaces.

mac — Specifies that the provisioned ieee-address will be used in the anti-spoofing entries for this SAP. The parameter may be changed mid-session. Existing sessions will not be re-programmed unless a **tools perform** command is issued for the lease. This parameter is only applicable for group interfaces.

neighbor-resolution

Syntax [**no**] **neighbor-resolution**

Context config>service>ies>if>ipv6>dhcp6-relay
config>service>vprn>if>ipv6>dhcp6-relay

Description This command enables neighbor resolution with DHCPv6 relay.

The **no** form of this command disables neighbor resolution.

remote-id

Syntax [**no**] **remote-id**

Context config>service>ies>if>ipv6>dhcp6>option
config>service>vprn>if>ipv6>dhcp6>option

Description This command enables the sending of remote ID option in the DHCPv6 relay packet.

The client DHCP Unique Identifier (DUID) is used as the remote ID.

The **no** form of this command disables remote ID.

remote-id

Syntax **remote-id** [**mac** | **string** *string*]
no remote-id

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option

Description This command specifies what information goes into the remote ID sub-option in the DHCP Relay packet.

If disabled, the **remote-id** sub-option of the DHCP packet is left empty.

The **no** form of this command returns the system to the default.

Parameters **mac** — Specifies the MAC address of the remote end is encoded in the sub-option

string — Specifies the remote ID, up to 32 characters

option

Syntax [**no**] **option**

Context config>service>ies>if>ipv6>dhcp6-relay
config>service>vprn>if>ipv6>dhcp6-relay

Description This command enables the context to configure DHCPv6 relay information options.

The **no** form of this command disables DHCPv6 relay information options.

interface-id

Syntax **interface-id** [**ascii-tuple**]
interface-id ifindex
interface-id sap-id
interface-id string *string*
no interface-id

Context config>service>ies>if>ipv6>dhcp6>option
config>service>ies>if>ipv6>dhcp6-relay
config>service>vprn>if>ipv6>dhcp6>option
config>service>vprn>if>ipv6>dhcp6-relay

Description This command enables the sending of interface ID options in the DHCPv6 relay packet.

The **no** form of this command disables the sending of interface ID options in the DHCPv6 relay packet

Parameters **ascii-tuple** — Specifies that the ASCII-encoded concatenated tuple is used which consists of the access-node-identifier, service-id, and interface-name, separated by "|"

ifindex — Specifies that the interface index is used. The If Index of a router interface can be displayed using the **show>router>interface>detail** command.

sap-id — Specifies that the SAP identifier is used

string — Specifies that a string is used

proxy-nd-policy

Syntax **proxy-nd-policy** *policy-name* [*policy-name*]
no proxy-nd-policy

Context config>service>ies>if>ipv6

Description This command configures a proxy neighbor discovery policy for the interface. This policy determines networks and sources for which proxy ND is attempted, when local proxy neighbor discovery is enabled.

The **no** form of this command reverts to the default value.

Parameters *policy-name* — Specifies up to five the export route policy names. Allowed values are any string, up to 32 characters, composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

The specified name(s) must already be defined.

option

Syntax [**no**] **option**

Context config>service>ies>if>dhcp
config>service>ies>sub-if>dhcp
config>service>ies>sub-if>grp-if>dhcp
config>service>vpls>sap>dhcp
config>service>vprn>if>dhcp
config>service>vprn>sub-if>dhcp
config>service>vprn>sub-if>grp-if>dhcp
config>subscr-mgmt>msap-policy>vpls-only>dhcp

Description This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options.

The **no** form of this command reverts to the default.

action

Syntax **action** {*dhcp-action*}
no action

Context config>service>ies>if>dhcp>option
config>service>ies>sub-if>grp-if>dhcp
config>service>vpls>sap>dhcp>option
config>service>vprn>if>dhcp>option
config>service>vprn>sub-if>grp-if>dhcp>option
config>subscr-mgmt>msap-policy>vpls-only>dhcp>option

Description This command configures the Relay Agent Information Option (Option 82) processing.

The **no** form of this command reverts to the default.

Default The default is to keep the existing information intact.

Parameters *dhcp-action* — Specifies the DHCP Relay reforwarding policy action.

Values replace — In the upstream direction (from the user), the Option 82 field from the router is inserted in the packet (overwriting any existing Option 82 field). In the downstream direction (toward the user) the Option 82 field is stripped (in accordance with RFC 3046).

drop — The DHCP packet is dropped if an Option 82 field is present, and a counter is incremented.

keep — The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is forwarded towards the client.

In Vendor-Specific Options (VSOs) scenarios, the behavior is slightly different. Even with the action=keep, the router will insert his own VSO into the Option 82 field. This will only be done when the incoming message has already an Option 82 field.

If no Option 82 field is present, the router will not create the Option 82 field, then, no VSO will be added to the message.

circuit-id

Syntax **circuit-id**
circuit-id {**ascii-tuple** | **if-index** | **sap-id** | **vlan-ascii-tuple**}
circuit-id hex [*hex-string*]
no circuit-id

Context config>service>ies>if>dhcp>option
config>service>ies>sub-if>grp-if>dhcp>option
config>service>vpls>sap>dhcp>option
config>service>vprn>if>dhcp>option

config>service>vprn>sub-if>grp-if>dhcp>option
config>subscr-mgmt>msap-policy>vpls-only>dhcp>option

Description When enabled, the router sends an ASCII-encoded tuple in the **circuit-id** sub-option of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by "|".

In order to send a tuple in the circuit ID, the **action replace** command must be configured in the same context.

If disabled, the **circuit-id** sub-option of the DHCP packet will be left empty.

The **no** form of this command reverts to the default.

Default circuit-id ascii-tuple

Parameters **ascii-tuple** — Specifies that the ASCII-encoded concatenated tuple consisting of the access-node-identifier, service-id, and interface-name is used

ifindex — Specifies that the interface index is used. The If Index of a router interface can be displayed using the command **show>router>if>detail**.

sap-id — Specifies that the SAP identifier is used

vlan-ascii-tuple — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and qinq ports only. Thus, when the Option 82 bits are stripped, dot1p bits will be copied to the Ethernet header of an outgoing packet.

hex-string — Specifies the hex value of this option

Values 0x0 to 0xFFFFFFFF...(up to 64 hex nibbles)

remote-id

Syntax **remote-id**
remote-id hex [*hex-string*]
remote-id mac
remote-id string *string*
no remote-id

Context config>service>ies>if>dhcp>option
config>service>ies>sub-if>grp-if>dhcp>option
config>service>vprn>sub-if>grp-if>dhcp>option
config>service>vpls>sap>dhcp>option
config>service>vprn>if>dhcp>option
config>subscr-mgmt>msap-policy>vpls-only>dhcp>option

Description This command specifies what information goes into the remote-id sub-option in the DHCP relay packet.

If disabled, the **remote-id** sub-option of the DHCP packet will be left empty.

The **no** form of this command reverts to the default.

Parameters **mac** — Specifies the MAC address of the remote end is encoded in the sub-option

string — Specifies the remote-id, up to 32 characters

hex-string — Specifies the hex value of this option

Values 0x0 to 0xFFFFFFFF...(up to 64 hex nibbles)

vendor-specific-option

Syntax [**no**] **vendor-specific-option**

Context config>service>vpls>sap>dhcp>option
config>service>ies>if>dhcp>option
config>service>vprn>if>dhcp>option
config>service>ies>sub-if>grp-if>dhcp>option
config>service>vprn>sub-if>grp-if>dhcp>option
config>subscr-mgmt>msap-policy>vpls-only>dhcp>option

Description This command enables the Nokia vendor-specific sub-option of the DHCP relay packet.

The **no** form of this command reverts to the default.

client-mac-address

Syntax [**no**] **client-mac-address**

Context config>service> ies>if>dhcp>option>vendor
config>service>vprn>if>dhcp>option>vendor
config>service>vprn>sub-if>grp-if>dhcp>option>vendor
config>service>ies>sub-if>grp-if>dhcp>option>vendor
config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor

Description This command enables the sending of the MAC address in the Nokia vendor-specific sub-option of the DHCP relay packet.

The **no** form of this command disables the sending of the MAC address in the Nokia vendor-specific sub-option of the DHCP relay packet.

sap-id

Syntax [**no**] **sap-id**

Context config>service>vpls>sap>dhcp>option>vendor

config>service>vprn>if>dhcp>option>vendor
config>service>vprn>sub-if>grp-if>dhcp>option>vendor
config>service>ies>sub-if>grp-if>dhcp>option>vendor
config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor

Description This command enables the sending of the SAP ID in the Nokia vendor-specific sub-option of the DHCP relay packet.

The **no** form of this command disables the sending of the SAP ID in the Nokia vendor-specific sub-option of the DHCP relay packet.

service-id

Syntax [**no**] **service-id**

Context config>service>vpls>sap>dhcp>option>vendor
config>service>vprn>if>dhcp>option>vendor
config>service>vprn>sub-if>grp-if>dhcp>option>vendor
config>service>ies>sub-if>grp-if>dhcp>option>vendor

Description This command enables the sending of the service ID in the Nokia vendor-specific sub-option of the DHCP relay packet.

The **no** form of this command disables the sending of the service ID in the Nokia vendor-specific sub-option of the DHCP relay packet.

string

Syntax [**no**] **string** *text*

Context config>service>vpls>sap>dhcp>option>vendor
config>service>vprn>if>dhcp>option>vendor
config>service>vprn>sub-if>grp-if>dhcp>option>vendor
config>service>ies>sub-if>grp-if>dhcp>option>vendor
config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor

Description This command specifies the string in the Nokia vendor-specific sub-option of the DHCP relay packet.

The **no** form of this command reverts to the default.

Parameters *text* — Specifies a string that can be any combination of ASCII characters, up to 32 characters. If spaces are used in the string, enclose the entire string in quotation marks (“ “).

system-id

Syntax [**no**] **system-id**

Context config>service>vpls>sap>dhcp>option>vendor
config>service>vprn>if>dhcp>option>vendor
config>service>vprn>sub-if>grp-if>dhcp>option>vendor
config>service>ies>sub-if>grp-if>dhcp>option>vendor
config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor

Description This command specifies whether the system-id is encoded in the Nokia vendor-specific sub-option of Option 82.

The **no** form of this command reverts to the default.

proxy-server

Syntax **proxy-server**

Context config>service>ies>if>dhcp
config>service>ies>sub-if>grp-if>dhcp
config>service>vpls>sap>dhcp
config>service>vprn>sub-if>grp-if>dhcp

Description This command enables the context to configure DHCP proxy server parameters.

emulated-server

Syntax **emulated-server** *ip-address*
no emulated-server

Context config>service>ies>if>dhcp>proxy-server
config>service>ies>sub-if>grp-if>dhcp>proxy-server
config>service>vpls>sap>dhcp>proxy-server
config>service>vprn>if>dhcp>proxy-server
config>service>ies>sub-if>grp-if>dhcp>proxy-server
config>service>vprn>sub-if>grp-if>dhcp>proxy-server
config>service>vprn>sub-if>ipv6>dhcp6>proxy
config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy-server

Description This command configures the IP address which will be used as the DHCP server address in the context of the SAP. Typically, the configured address should be in the context of the subnet represented by the service.

The **no** form of this command reverts to the default setting. The local proxy server will not become operational without the emulated-server address being specified.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the emulated server's IP address. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

lease-time

Syntax **lease-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*] [**override**]
no lease-time

Context config>service>vpls>sap>dhcp>proxy-server
config>service>ies>if>dhcp>proxy-server
config>service>ies>sub-if>grp-if>dhcp>proxy-server
config>service>vprn>if>dhcp>proxy
config>service>vprn>sub-if>grp-if>dhcp>proxy-server
config>service>ies>sub-if>grp-if>wpp
config>service>vprn>sub-if>grp-if>wpp

Description This command defines the length of lease-time that will be provided to DHCP clients. By default the local-proxy-server will always make use of the lease-time information provide by either a RADIUS or DHCP server.

The **no** form of this command disables the use of the lease-time command. The local-proxy-server will use the lease-time offered by either a RADIUS or DHCP server.

Default 7 days 0 hours 0 seconds

Parameters **override** — Specifies that the local-proxy-server will use the configured lease-time information to provide DHCP clients

days — Specifies the number of days that the given IP address is valid

Values 0 to 3650

hours — Specifies the number of hours that the given IP address is valid

Values 0 to 23

minutes — Specifies the number of minutes that the given IP address is valid

Values 0 to 59

seconds — Specifies the number of seconds that the given IP address is valid

Values 0 to 59

pool-name

Syntax [**no**] **pool-name**

Context config>service>vprn>if>dhcp>option>vendor
config>service>vprn>sub-if>grp-if>dhcp>option>vendor

config>service>ies>if>dhcp>option>vendor
config>service>ies>sub-if>grp-if>dhcp>option>vendor

Description This command sends the pool name in the Nokia vendor specific sub-option of the DHCP relay packet.

The **no** form of this command reverts to the default.

override-slaac

Syntax [**no**] **override-slaac**

Context config>service>vprn>sub-if>ipv6>dhcp6
config>service>ies>sub-if>ipv6>dhcp6
config>service>vprn>sub-if>grp-if>ipv6>dhcp6
config>service>ies>sub-if>grp-if>ipv6>dhcp6

Description This command allows a DHCP IA_NA address to override and replace a host existing SLAAC address. When this feature is enabled, a subscriber SLAAC address is removed once the DHCP IA_NA address assignment is completed. If used with conjunction with the **allow-multiple-wan-address** command, the DHCP IA_NA address will also override the SLAAC address.

pd-managed-route

Syntax **pd-managed-route** [**next-hop** {**ipv4** | **ipv6**}]
no pd-managed-route

Context config>service>vprn>sub-if>ipv6>dhcp6
config>service>ies>sub-if>ipv6>dhcp6
config>service>vprn>sub-if>grp-if>ipv6>dhcp6
config>service>ies>sub-if>grp-if>ipv6>dhcp6

Description This command enables DHCP IA-PD (delegated prefix) to be modeled as managed (framed) route instead of as a subscriber-host. Antispoof filtering for the subscriber host associated with the IA-PD route must be set to **nh-mac**. The subscriber specific parameters (such as **sla-profile** or **sub-profile**) will be ignored during the authentication phase because IA-PD is not modeled as a subscriber host. Other subscriber host-specific functions (for example, host overrides via CoA or host accounting) are not possible with a PD as the managed route.

By default, or when configured with the **next-hop ipv6** parameter, the next-hop for PD managed route is an IPv6 WAN sub-host (DHCP IA-NA or SLAAC) with the same mac address as the one in the DHCP lease state for the managed IA-PD. The DHCP IA-NA next-hop host will always override the SLAAC next-hop host if both are available. If the IPv6 next-hop is not present when the framed IA-PD is instantiated, the IA-PD will be set up but the PD managed route will not be installed in the IPv6 route table and the DHCPv6 lease state for the IA-PD will have the managed route status (DHCP6 MRt Status) set to “noNextHop”.

When configured with the **next-hop ipv4** parameter the next-hop for PD managed route is a DHCPv4 sub-host that belongs to the same IPoE session or PPPoE session. For IPoE, **ipoe-session** must be enabled on the group-interface. If **ipoe-session** is disabled, an IPv4 next-hop will not be found. If the IPv4 next-hop is not found or not present at the time when the framed IA-PD is instantiated, the IA-PD will be set up but the PD managed route will not be installed in the IPv6 route table. In this case, the DHCPv6 lease state for the IA-PD will have the managed route status (DHCP6 MRt Status) set to noNextHop.

Note: IPv6 filters, QoS IPv6 criteria, and IPv6 multicast are not supported for DHCPv6 IA-PD as managed route pointing to an IPv4 subscriber host as next-hop.

The DHCP IA-PD modeled as a route is displayed differently than regular subscriber hosts in **show** commands related to subscriber host state. The PD managed route is always shown directly below the host it is using as the next hop. The forwarding status of the PD managed route is also shown, where (N) indicate that the PD managed route is not forwarding. In addition, DHCP IA-PD route is displayed as a managed route for the corresponding IPv6 subscriber host (DHCP IA-NA or SLAAC) or DHCPv4 subscriber host.

DHCP IA-PD information for managed IA-PD route is still maintained in the DHCPv6 lease state.

The **no** form of this command reverts to the default.

Parameters **next-hop** — Specifies the next-hop type for the DHCP IA-PD managed route

Values **ipv4** - The next-hop for PD managed route is a DHCPv4 sub-host that belongs to the same IPoE session (based on the IPoE session key which is **sap-mac** by default). IPoE session must be enabled on the group-interface.

ipv6 - The next-hop for PD managed route is an IPv6 WAN sub-host (DHCP IA-NA or SLAAC) with the same MAC address as the one in the DHCP lease state for the managed IA-PD. This is the default when no next-hop is specified.

enable-ingress-stats

Syntax [**no**] **enable-ingress-stats**

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command enables the collection of ingress interface IP stats. This command is only applicable to IP statistics, and not to uRPF statistics.

If enabled, then the following statistics are collected:

- IPv4 offered packets
- IPv4 offered octets
- IPv6 offered packets
- IPv6 offered octets

Octet statistics for IPv4 and IPv6 bytes at IP interfaces include the Layer 2 frame overhead.

The **no** form of the reverts to the default.

duid

Syntax **duid** *duid* [**iaid** *iaid*]
no duid

Context config>service>ies>if>ipv6>dhcp6>pfx-delegate>prefix

Description This command configures the DHCP Unique Identifier (DUID) of the DHCP client.

The **no** form of this command reverts to the default.

Default duid 2

Parameters *duid* — Specifies the ID of the requesting router, up to a maximum of 128 hex values. If set to a non-zero value the prefix defined will only be delegated to this router. If set to zero, the prefix will be delegated to any requesting router.

iaid — Specifies the identity association identification (IAID) from the requesting router that needs to match in order to delegate the prefix defined in this row. If set to 0 no match on the received IAID is done.

Values 1 to 4294967295

preferred-lifetime

Syntax **preferred-lifetime** *seconds*
preferred-lifetime infinite
no preferred-lifetime

Context config>service>ies>if>ipv6>dhcp6>pfx-delegate>prefix
config>service>vprn>if>ipv6>dhcp6>pfx-delegate>prefix
config>service>vprn>sub-if>ipv6>dhcp6>proxy
config>service>ies>sub-if>ipv6>dhcp6>proxy
config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy
config>service>ies>sub-if>grp-if>ipv6>dhcp6>proxy

Description This command configures the IPv6 prefix/mask preferred life time. The preferred-lifetime value cannot be bigger than the valid-lifetime value.

The **no** form of this command reverts to the default value.

Default preferred-lifetime 604800 (7 days)

Parameters *seconds* — Specifies the preferred lifetime in seconds

Values 1 to 4294967294

infinite — Specifies that the remaining time will never expire. The value 4294967294 is interpreted as infinite.

preferred-lifetime

Syntax **preferred-lifetime** *seconds*
preferred-lifetime infinite
no preferred-lifetime

Context config>service>ies>sub-if>grp-if>ipv6>rtr-adv>pfx-opt
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv>pfx-opt
config>service>ies>sub-if>ipv6>rtr-adv>pfx-opt
config>service>vprn>sub-if>ipv6>rtr-adv>pfx-opt
config>router>dhcp6>server>pool>prefix>options
config>service>vprn>dhcp6>server>pool>prefix>options

Description This command specifies the remaining time for this prefix to be preferred, thus time until deprecation.

The **no** form of this command reverts to the default.

Default preferred-lifetime 3600

Parameters *seconds* — Specifies the time for the prefix to remain preferred on this group-interface in seconds

Values 0 to 4294967295

infinite — Specifies that the remaining time will never expire. The value 4294967295 is interpreted as infinite.

valid-lifetime

Syntax **valid-lifetime** *seconds*
valid-lifetime infinite
no valid-lifetime

Context config>service>ies>if>ipv6>dhcp6>pfx-delegate>prefix
config>service>vprn>sub-if>grp-if>ipv6>dhcp6
config>service>ies>sub-if>grp-if>ipv6>dhcp6

Description This command configures the time, in seconds, that the prefix is valid.

The **no** form of this command reverts to the default value.

Default valid-lifetime 2592000 (30 days)

Parameters *seconds* — Specifies the time, in seconds, that this prefix remains valid

Values 1 to 4294967294

infinite — Specifies that this prefix remains valid infinitely. The value 4294967295 is interpreted as infinite.

valid-lifetime

Syntax **valid-lifetime** [*valid-lifetime*]

Context config>service>vprn>dhcp6>local-dhcp-server>pool>prefix
config>router>dhcp6>server>pool>prefix

Description This command configures the valid lifetime for the IPv6 prefix or address in the option.

The **no** form of this command reverts to the default.

Default valid-lifetime days 1

Parameters *valid-lifetime* — Specifies the valid lifetime

Values

days *days*	0 to 3650
hrs *hours*	0 to 23
min *minutes*	0 to 59
sec *seconds*	0 to 5

python-policy

Syntax **python-policy** *name*
no python-policy

Context config>service>ies>if>dhcp
config>service>vprn>if>dhcp
config>service>ies>sub-if>dhcp
config>service>vprn>sub-if>dhcp

Description This command specifies the python-policy to be used for DHCPv4 relay.

The **no** form of this command reverts to the default.

Parameters *name* — Specifies the name of an existing python script, up to 32 characters.

python-policy

Syntax **python-policy** *name*
no python-policy

Context config>service>vprn>sub-if
config>service>vprn>sub-if>ipv6>dhcp6
config>service>ies>sub-if>ipv6>dhcp6
config>service>ies>if>ipv6>dhcp6-relay
config>service>vprn>if>ipv6>dhcp6-relay
config>service>ies>sub-if>dhcp
config>service>vprn>sub-if>dhcp
config>service>ies>sub-if>grp-if>dhcp
config>service>vprn>sub-if>grp-if>dhcp
config>service>ies>sub-if>grp-if>pppoe
config>service>vprn>sub-if>grp-if>pppoe
config>service>vprn>sub-if>ipv6>dhcp6>proxy
config>service>vprn>sub-if>grp-if>ipv6>dhcp6>proxy-server

Description This command specifies the python-policy to be used for DHCPv6 relay.

The **no** form of this command reverts to the default.

Parameters *name* — Specifies the name of an existing python script, up to 32 characters

relay

Syntax [**no**] **relay**

Context config>service>vprn>sub-if>grp-if>ipv6>dhcp6
config>service>ies>sub-if>grp-if>ipv6>dhcp6
config>service>vprn>sub-if>ipv6>dhcp6
config>service>ies>sub-if>ipv6>dhcp6

Description This command enables the context to configure DHCPv6 relay parameters for this interface.

snoop

Syntax [**no**] **snoop**

Context config>service>vpls>sap>dhcp
config>service>vpls>spoke-sdp>dhcp
config>service>vpls>mesh-sdp>dhcp
config>service>vprn>if>dhcp>option

Description This command enables DHCP snooping of DHCP messages on the SAP or SDP. Enabling DHCP snooping on interfaces (SAPs and SDP bindings) is required where DHCP messages important to lease state table population are received, or where Option 82 information is to be inserted. This includes interfaces that are in the path to receive messages from either DHCP servers or from subscribers.

The **no** form of this command disables DHCP snooping on the specified SAP or SDP binding.

dhcp-user-db

Syntax **dhcp-user-db** *local-user-db-name*
no dhcp-user-db

Context config>service>vpls>sap

Description This command enabled access to LUDB for DHCPv4 hosts under the capture SAP. The name of this local user database must match the name of local user database configured under the **config>service>vprn/ies>sub-if>group-if>dhcp** hierarchy.

Parameters *local-user-db* — Specifies the name of the local-user-database, up to 32 characters

dhcp-python-policy

Syntax **dhcp-python-policy** *policy-name*
no dhcp-python-policy

Context config>service>vpls>sap

Description This command specifies the name of the Python policy. The Python policy is created in the **config>python>python-policy** *name* context.

The **no** form of this command reverts to the default.

Parameters *policy-name* — Specifies a Python policy name, up to 32 characters

dhcp6

Syntax **dhcp6**

Context config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>grp-if>ipv6

Description This command enables the context to configure DHCP6 parameters.

interface-id

Syntax **interface-id**
interface-id ascii-tuple
interface-id vlan-ascii-tuple
no interface-id

Context config>service>ies>if>ipv6>dhcp6>option
config>service>vprn>if>ipv6>dhcp6>option

Description This command configure the interface-id sub-option of the DHCP6 Relay packet

The **no** form of this command disables the sending of interface ID options in the DHCPv6 relay packet.

Parameters **ascii-tuple** — Specifies that the ASCII-encoded concatenated tuple will be used which consists of the access-node-identifier, service-id, and interface-name, separated by "|"

vlan-ascii-tuple — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already

mac — Specifies the MAC address of the remote end is encoded in the sub-option

remote-id

Syntax **remote-id**
remote-id mac
remote-id string [*32 chars max*]
no remote-id

Context config>service>ies>if>ipv6>dhcp6>option
config>service>vprn>if>ipv6>dhcp6>option

Description This command enables the sending of remote ID option in the DHCPv6 relay packet.

The client DHCP Unique Identifier (DUID) is used as the remote ID.

The **no** form of this command disables the sending of remote ID option in the DHCPv6 relay packet.

dhcp6-user-db

Syntax **dhcp6-user-db** *local-user-db*
no dhcp6-user-db

Context config>service>vpls>sap

Description This command enabled access to LUDB for DHCPv6 hosts under the capture SAP. The name of this LUDB must match the name of the LUDB configured under the **config>service>vprn/ies>sub-if>grp-if>dhcp** hierarchy.

The **no** form of this command reverts to the default.

Parameters *local-user-db* — Specifies the name of the local-user-database, up to 32 characters

ppp-user-db

Syntax **ppp-user-db** *local-user-db-name*
no ppp-user-db

Context config>service>vpls

Description This command enabled access to LUDB for PPPoE and PPPoEoA v4and v6 hosts under the capture SAP. The name of this local user database must match the name of local user database configured under the **config>service>vprn/ies>sub-if>grp-if>pppoe** hierarchy.

The **no** form of this command reverts to the default.

Parameters *local-user-db* — Specifies the name of the local-user-database, up to 256 characters

pppoe-user-db

Syntax **pppoe-user-db** *local-user-db-name*
no pppoe-user-db

Context config>service>vpls

Description This command enabled access to LUDB for PPPoE and PPPoEoA v4and v6 hosts under the capture SAP. The name of this LUDB must match the name of the LUDB configured under the **config>service>vprn/ies>sub-if>grp-if>pppoe** hierarchy.

The **no** form of this command reverts to the default.

Parameters *local-user-db* — Specifies the name of the local user database, up to 256 characters

filter

Syntax **filter** *filter-id*
no filter

Context config>service>ies>sub-if>grp-if>dhcp
config>service>vprn>sub-if>grp-if>dhcp

Description This command assigns a DHCP filter to the group-interface. This feature is used where the SR 7750 is the second DHCP relay or where DHCP messages are snooped for subscriber management. The filter can be used to bypass host creation, drop DHCP message, or perform no action.

The **no** form of this command reverts to the default.

Parameters *filter-id* — Specifies the DHCP filter ID for this interface

Values 1 to 65535

filter

Syntax **filter** *filter-id*
no filter

Context config>service>ies>sub-if>grp-if>dhcp6
config>service>vprn>sub-if>grp-if>dhcp6

Description This command assigns a DHCP6 filter to the group interface. This feature is used where the SR 7750 is the second DHCP6 relay or where DHCP6 messages are snooped for subscriber management. The filter can be used to bypass host creation, drop DHCP6 message, or perform no action.

The **no** form of this command reverts to the default.

Parameters *filter-id* — Specifies the DHCP6 filter ID for this interface

Values 1 to 65535

gi-address

Syntax **gi-address** *ip-address* [**src-ip-address**]
no gi-address

Context config>service>ies>if>dhcp
config>service>vprn>if>dhcp
config>service>vprn>sub-if>dhcp
config>service>ies>sub-if>grp-if>dhcp
config>service>ies>sub-if>dhcp

Description This command configures the gateway interface address for the DHCP relay. A subscriber interface can include multiple group interfaces with multiple SAPs. The GI address is needed, when the router functions as a DHCP relay, to distinguish between the different subscriber interfaces and potentially between the group interfaces defined.

By default, the GI address used in the relayed DHCP packet is the primary IP address of a normal IES interface. Specifying the GI address allows the user to choose a secondary address. For group interfaces a GI address must be specified under the group interface DHCP context or subscriber-interface DHCP context in order for DHCP to function.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the host IP address to be used for DHCP relay packets.

src-ip-address — Specifies that this GI address is to be the source IP address for DHCP relay packets

gi-address

Syntax **gi-address** *ip-address*
no gi-address

Context config>subscr-mgmt>loc-user-db>ipoe>host

Description This command allows selection of gi-addresses based on the host entry in LUDB.

The gi-address must be a valid address (associated with an interface) within the routing context that received the DHCP message on the access side.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the IPv4 GI address

relay-plain-bootp

Syntax [**no**] **relay-plain-bootp**

Context config>service>ies>if>dhcp
config>service>vprn>if>dhcp

Description This command enables the relaying of plain BOOTP packets.

The **no** form of this command disables the relaying of plain BOOTP packets.

relay-proxy

Syntax **relay-proxy** [**release-update-src-ip**] [**siaddr-override** *ip-address*]
no relay-proxy

Context config>service>ies>if>dhcp
config>service>ies>sub-if>dhcp
config>service>ies>sub-if>grp-if>dhcp

config>service>vprn>if>dhcp
config>service>vprn>sub-if>dhcp
config>service>vprn>sub-if>grp-if>dhcp

Description This command enables the DHCPv4 relay proxy function on the interface. The command has no effect when no dhcp servers are configured (DHCPv4 relay not configured). By default, unicast DHCPv4 release messages are forwarded transparently.

A relay proxy enhances the relay such that it also relays unicast client DHCPv4 REQUEST messages (lease renewals).

- In the upstream direction, update the source IP address and add the gateway IP address (gi-address) field before sending the message to the intended DHCP server (the message is not broadcasted to all configured DHCP servers.
- In the downstream direction, remove the gi-address and update the destination IP address to the address of the `yiaddr` (your IP address) field.

The optional **release-update-src-ip** parameter updates the source IP address of a DHCP RELEASE message with the address used for relayed DHCPv4 messages.

The optional **siaddr-override** *ip-address* parameter enables DHCP server IP address hiding towards the client. This parameter requires that **lease-populate** is enabled on the interface. The DHCP server ip address is required for the address hiding function and is stored in the lease state record. The client interacts with the relay proxy as if it is the DHCP server. In all DHCP messages to the client, the value of following header fields and DHCP options containing the DHCP server IP address is replaced with the configured <*ip-address*>:

- the “source IP address” field in the IP DHCPv4 packet header
- the “siaddr” field in the DHCPv4 header if not equal to zero in the message received from the server
- the Server Identification option (DHCPv4 option 54) if present in the original server message
- the source IP address field in the IP packet header

DHCP OFFER selection during initial binding is done in the relay-proxy. Only the first DHCP OFFER message is forwarded to the client. Subsequent DHCP OFFER messages from different servers are silently dropped.

Parameters **release-update-src-ip** — Updates the source IP address of a DHCP RELEASE message with the address used for relayed DHCPv4 messages.

ip-address — Enables DHCPv4 server address hiding towards the DHCPv4 client and activates DHCPv4 OFFER selection in case multiple DHCP servers are configured. The *ip-address* can be any local address in the same routing instance. If DHCP relay lease-split is enabled, **siaddr-override** *ip-address* has priority over the **emulated-server** *ip-address* configured in the proxy-server and will be used as the source IP address.

server

Syntax **server** *ipv6z-address* [*ipv6z-address*]
no server [*ipv6z-address*]

Context config>service>ies>if>ipv6>dhcp6

Description This command specifies a list of servers where DHCP6 requests will be forwarded. The list of servers can entered as either IP addresses or fully qualified domain names. There must be at least one server specified for DHCP6 relay to work. If there are multiple servers then the request is forwarded to all of the servers in the list.

The **no** form of this command reverts to the default.

Parameters *ipv6z-address* — Specifies up to eight non-global IPv4 addresses including a zone index as defined by the InetAddressIPv4z textual convention

Values

ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D

server

Syntax **server** *server1* [*server2*]

Context config>service>ies>if>dhcp
config>service>vprn>if>dhcp
config>service>ies>sub-if>grp-if>dhcp

Description This command specifies a list of servers where requests will be forwarded. The list of servers can entered as either IP addresses or fully qualified domain names. There must be at least one server specified for DHCP relay to work. If there are multiple servers then the request is forwarded to all of the servers in the list.

There can be a maximum of 8 DHCP servers configured.

The **no** form of this command reverts to the default.

Parameters *server* — Specifies, up to eight DHCP server IP addresses

virtual-subnet

Syntax [**no**] **virtual-subnet**

Context config>service>ies>sub-if>dhcp
config>service>vprn>sub-if>dhcp

Description This command enables a virtual-subnet for DHCPv4 hosts under the subscriber-interface. With this command configured, the system will snoop and record the default router address in the DHCP ACK message for the DHCPv4 ESM host. The system could answer ping or traceroute request even if the default router address is not configured on the subscriber-interface.

The **no** form of this command reverts to the default.

relay-plain-bootp

Syntax [**no**] **relay-plain-bootp**

Context config>service>vprn>if>dhcp

Description This command enables the relaying of plain BOOTP packets.

The **no** form of this command disables the relaying of plain BOOTP packets.

use-arp

Syntax [**no**] **use-arp**

Context config>service>vprn>if>dhcp
config>service>ies>if>dhcp

Description This command enables the use of ARP to determine the destination hardware address.

The **no** form of this command disables the use of ARP to determine the destination hardware address.

user-db

Syntax **user-db** *local-user-db-name*
no user-db

Context config>service>ies>sub-if>grp-if>pppoe
config>service>vprn>sub-if>grp-if>pppoe
config>service>ies>sub-if>grp-if>ppp
config>service>vprn>sub-if>grp-if>ppp

Description This command configures the local user database to use for PPP PAP/CHAP authentication

The **no** form of this command reverts to the default.

Parameters *local-user-db-name* — Specifies the local user database name, up to 32 characters

trusted

Syntax [**no**] **trusted**

Context config>service>ies>if>dhcp
config>service>vprn>if>dhcp
config>service>ies>sub-if>grp-if>dhcp

Description This command enables relaying untrusted packets. According to RFC 3046, *DHCP Relay Agent Information Option*, a DHCP request where the giaddr is 0.0.0.0 and which contains an Option 82 field in the packet, should be discarded, unless it arrives on a “trusted” circuit. If trusted mode is enabled on an IP interface, the Relay Agent (the router) modifies the request's giaddr to be equal to the ingress interface and forward the request.

Note: This behavior only applies when the action in the Relay Agent Information Option is **keep**. In the case where the option 82 field is being replaced by the Relay Agent (action = **replace**), the original Option 82 information is lost anyway, and thus, there is no reason for enabling the trusted option.

The **no** form of this command reverts to the default.

ipv6

Syntax **ipv6**

Context config>service>ies>if>
config>service>vprn>if>
config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command enables the context to configure IPv6 parameters for the interface.

address

Syntax **address** *ipv6-address/prefix-length* [**eui-64**]
no address *ipv6-address/prefix-length*

Context config>service>ies>if>ipv6

Description This command assigns an IPv6 address to the IES interface.

The **no** form of this command reverts to the default.

Parameters *ipv6-address/prefix-length* — Specifies the IPv6 address on the interface

Values

ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d x [0 to FFFF]H d [0 to 255]D
prefix-length	1 to 128

eui-64 — When the **eui-64** keyword is specified, a complete IPv6 address from the supplied prefix and 64-bit interface identifier is formed. The 64-bit interface identifier is derived from MAC address on Ethernet interfaces. For interfaces without a MAC address, for example ATM interfaces, the Base MAC address of the chassis is used.

allow-multiple-wan-addresses

Syntax [**no**] **allow-multiple-wan-addresses**

Context config>service>ies>sub-if>ipv6
config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6
config>service>vprn>sub-if>grp-if>ipv6

Description This command enables host to have two WAN addresses, one from DHCP IA_NA and one from SLAAC assignment.

The **no** form of this command reverts to the default.

dhcp6-relay

Syntax [**no**] **dhcp6-relay**

Context config>service>vprn>if>ipv6

Description This command enables the context to configure DHCPv6 relay parameters for the interface.

dhcp

Syntax **dhcp**

Context config>service>vprn>if
config>service>vprn>
config>service>vprn>sub-if>grp-if

Description This command enables the context to configure DHCP parameters.

action

Syntax **action** {**replace** | **drop** | **keep**}
no action

Context config>service>vprn>if>dhcp>option
config>service>vprn>sub-if>grp-if>dhcp>option

Description This command configures the processing required when the SR-Series receives a DHCP request that already has a Relay Agent Information Option (Option 82) field in the packet.

The **no** form of this command returns the system to the default value.

Default action keep — Per RFC 3046, *DHCP Relay Agent Information Option*, section 2.1.1, *Reforwarded DHCP requests*. The default is to keep the existing information intact. The exception to this is if the giaddr of the received packet is the same as the ingress address on the router. In that case the packet is dropped and an error is logged.

Parameters **replace** — In the upstream direction (from the user), the existing Option 82 field is replaced with the Option 82 field from the router. In the downstream direction (towards the user) the Option 82 field is stripped (in accordance with RFC 3046).

drop — The packet is dropped, and an error is logged

keep — The existing information is kept in the packet and the router does not add any additional information. In the downstream direction the Option 82 field is not stripped and is sent on towards the client.

The behavior is slightly different in case of Vendor Specific Options (VSOs). When the keep parameter is specified, the router will insert his own VSO into the Option 82 field. This will only be done when the incoming message has already an Option 82 field.

If no Option 82 field is present, the router will not create the Option 82 field. In this in that case, no VSO will be added to the message.

match-circuit-id

Syntax [**no**] **match-circuit-id**

Context config>service>vprn>sub-if>grp-if>dhcp
config>service>ies>sub-if>grp-if>dhcp

Description This command enables Option 82 circuit ID on relayed DHCP packet matching. For routed CO, the group interface DHCP relay process is stateful. When packets are relayed to the server the virtual router ID, transaction ID, SAP ID, and client hardware MAC address of the relayed packet are tracked.

When a response is received from the server the virtual router ID, transaction ID, and client hardware MAC address must be matched to determine the SAP on which to send the packet out. In some cases, the virtual router ID, transaction ID, and client hardware MAC address are not guaranteed to be unique.

When the **match-circuit-id** command is enabled this part of the key is used to guarantee correctness in our lookup. This is really only needed when dealing with an IP aware DSLAM that proxies the client hardware MAC address.

The **no** form of this command reverts to the default.

option

Syntax [**no**] **option**

Context config>service>vprn>if>dhcp
config>service>vprn>sub-if>dhcp
config>service>vprn>sub-if>grp-if>dhcp
config>service>vpls>sap>dhcp
config>service>ies>if>dhcp
config>service>ies>sub-if>grp-if>dhcp

Description This command enables DHCP Option 82 (Relay Agent Information Option) parameters processing and enters the context for configuring Option 82 sub-options.

The **no** form of this command returns the system to the default.

vendor-specific-option

Syntax [**no**] **vendor-specific-option**

Context config>service>ies>sub-if>dhcp
config>service>vprn>if>dhcp>option
config>service>vprn>sub-if>grp-if>dhcp>option

Description This command configures the Nokia vendor specific sub-option of the DHCP relay packet.

The **no** form of this command reverts to the default.

client-mac-address

Syntax [**no**] **client-mac-address**

Context config>service>ies>sub-if>dhcp>option
config>service>vprn>if>dhcp>option>vendor
config>service>vprn>if>dhcp>option>vendor
config>service>vprn>sub-if>grp-if>dhcp>option>vendor

Description This command enables the sending of the MAC address in the Nokia vendor specific sub-option of the DHCP relay packet.

The **no** form of this command disables the sending of the MAC address in the Nokia vendor specific sub-option of the DHCP relay packet.

The **no** form of this command reverts to the default.

sap-id

Syntax [**no**] **sap-id**

Context config>service>vprn>if>dhcp>option>vendor
config>service>vprn>sub-if>grp-if>dhcp>option>vendor

Description This command enables the sending of the SAP ID in the Nokia vendor specific sub-option of the DHCP relay packet.

The **no** form of this command disables the sending of the SAP ID in the Nokia vendor specific sub-option of the DHCP relay packet.

service-id

Syntax [**no**] **service-id**

Context config>service>vprn>if>dhcp>option>vendor
config>service>vprn>sub-if>grp-if>dhcp>option>vendor

Description This command enables the sending of the service ID in the Nokia vendor specific sub-option of the DHCP relay packet.

The **no** form of this command disables the sending of the service ID in the Nokia vendor specific sub-option of the DHCP relay packet.

string

Syntax [**no**] **string** *text*

Context config>service>vprn>if>dhcp>option>vendor
config>service>vprn>sub-if>grp-if>dhcp>option>vendor

Description This command specifies the vendor specific sub-option string of the DHCP relay packet.

The **no** form of this command returns the default value.

Parameters *text* — The string can be any combination of ASCII characters, up to 32 characters. If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

system-id

Syntax [**no**] **system-id**

Context config>service>vprn>if>dhcp>option>vendor
config>service>vprn>sub-if>grp-if>dhcp>option>vendor

Description This command specifies whether the system-id is encoded in the Nokia vendor specific sub-option of Option 82.

The **no** form of this command reverts to the default.

host-connectivity-verify

Syntax **host-connectivity-verify source-ip** *ip-address* [**source-mac** *ieee-address*] [**interval** *interval*] [**action** {**remove** | **alarm**}] [**timeout** *retry-timeout*] [**retry-count** *count*]

Context config>service>vpls
config>service>vpls>sap

Description This command enables subscriber host connectivity verification on a given VPLS SAP or within a VPLS service.

This tool will periodically scan all known hosts (from dhcp-state) and perform a UC ARP request. The subscriber host connectivity verification will maintain state (connected vs. not-connected) for all hosts.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies an unused IP address in the same network for generation of subscriber host connectivity verification packets

ieee-address — Specifies the source MAC address to be used for generation of subscriber host connectivity verification packets

interval — The interval, in minutes, which specifies the time interval in which all known sources should be verified. The actual rate is then dependent on number of known hosts and interval.

Values 1 to 6000
Note that a zero value can be used by the SNMP agent to disable host-connectivity-verify.

action {**remove** | **alarm**} — Defines the action taken on a subscriber host connectivity verification failure for a given host. The **remove** keyword raises an alarm and removes dhcp-state and releases all allocated resources (queues, table entries, and so on). DHCP release will be signaled to corresponding DHCP server. Static host will be never removed. The **alarm** keyword raises an alarm indicating that the host is disconnected.

retry-timeout — Specifies the time, in seconds, to wait before a SAP that has been disabled after exceeding the maximum relearn rate is re-enabled

Values 0 to 120

count — Specifies the number of connectivity check retransmissions

Values 10 to 60 seconds

source-address

Syntax **source-address** *ipv6-address*
no source-address

Context config>service>ies>if>ipv6>dhcp6-relay
config>service>ies>sub-if>ipv6>dhcp6>relay
config>service>ies>sub-if>grp-if>ipv6>dhcp6>relay
config>service>vprn>if>ipv6>dhcp6-relay
config>service>vprn>sub-if>ipv6>dhcp6>relay
config>service>vprn>sub-if>grp-if>ipv6>dhcp6>relay

Description This command configures the source IPv6 address of the DHCPv6 relay messages.

The **no** form of this command reverts to the default.

Parameters *ipv6-address* — Specifies the source IPv6 address of the DHCPv6 relay messages

Values

ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x: [0 to FFFF]H
	d: [0 to 255]D

ipoe-bridged-mode

Syntax [**no**] **ipoe-bridged-mode**

Context config>service>ies>sub-if>ipv6
config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6
config>service>vprn>sub-if>grp-if>ipv6

Description This command enables IPv6 IPoE bridged mode.

The **no** form of this command disables the IPv6 IPoE bridged mode.

user-db

Syntax **user-db** *local-user-db-name*
no user-db

Context config>service>ies>if>ipv6>dhcp6-relay
config>service>vprn>if>ipv6>dhcp6-relay

Description This command enables access to the LUDB for DHCPv6 messages under a routed interface. The name of this LUDB must match the name of the LUDB configured by the **config>sub-gmt>local-user-db** command.

The **no** form of this command reverts to the default.

Parameters *local-user-db-name* — Specifies the name of the local user database, up to 32 characters

dhcp6-server

Syntax [**no**] **dhcp6-server**

Context config>service>ies>if>ipv6

Description This command enables the context to configure DHCPv6 server parameters for the IES interface.

The **no** form of this command disables the DHCP6 server.

max-nbr-of-leases

Syntax **max-nbr-of-leases** *max-nbr-of-leases*
no max-nbr-of-leases

Context config>service>ies>if>ipv6>dhcp6-server

Description This command configures the maximum number of lease states installed by the DHCP6 server function allowed on this interface.

The **no** form of this command to the default.

Default max-nbr-of-leases 8000

Parameters *max-nbr-of-leases* — Specifies the maximum number of lease states installed by the DHCP6 server function allowed on this interface

Values 0 to 8000

prefix-delegation

Syntax [**no**] **prefix-delegation**

Context config>service>ies>if>ipv6>dhcp6-server

Description This command enables the prefix delegation options for delegating a long-lived prefix from a delegating router to a requesting router, where the delegating router does not require knowledge about the topology of the links in the network to which the prefixes will be assigned.

The **no** form of this command disables prefix-delegation.

prefix

Syntax [**no**] **prefix** *ipv6-address/prefix-length*

Context config>service>ies>if>ipv6>dhcp6-server>pfx-delegate

Description This command specifies the IPv6 prefix that will be delegated by this system.

The **no** form of this command reverts to the default.

Parameters *ipv6-address/prefix-length* — Specifies the IPv6 address on the interface

Values

ipv6-address/prefix:	ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x [0 to FFFF]H
		d [0 to 255]D
	prefix-length	1 to 128

local-proxy-arp

Syntax [**no**] **local-proxy-arp**

Context config>service>vprn>if
config>service>vprn>sub-if>grp-if

Description This command enables local proxy ARP. When local proxy ARP is enabled on an IP interface, the system responds to all ARP requests for IP addresses belonging to the subnet with its own MAC address, and thus will become the forwarding point for all traffic between hosts in that subnet. When local-proxy-arp is enabled, ICMP redirects on the ports associated with the service are automatically blocked.

The **no** form of this command reverts to the default.

mac

Syntax [**no**] **mac** *ieee-mac-address*

Context config>service>vprn>if
config>service>vprn>if>vrrp
config>service>vprn>sub-if>grp-if

Description This command assigns a specific MAC address to a VPRN IP interface.

The **no** form of this command returns the MAC address of the IP interface to the default value.

Default The physical MAC address associated with the Ethernet interface that the SAP is configured on.

Parameters *ieee-mac-address* — Specifies the 48-bit MAC address for the static ARP in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

proxy-arp-policy

Syntax [**no**] **proxy-arp-policy** *policy-name* [*policy-name*)

Context config>service>vprn>if
config>service>vprn>sub-if>grp-if

Description This command enables a proxy ARP policy for the interface.

The **no** form of this command disables the proxy ARP capability.

Parameters *policy-name* — Specifies up to five export route policy names. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

The specified name(s) must already be defined.

redundant-interface

Syntax **redundant-interface** *red-ip-int-name*
no redundant-interface

Context config>service>vprn
config>service>vprn>sub-if>grp-if

Description This command configures a redundant interface used for dual homing.

The **no** form of this command reverts to the default.

Parameters *red-ip-int-name* — Specifies the redundant IP interface name

remote-proxy-arp

Syntax [**no**] **remote-proxy-arp**

Context config>service>vprn>if
config>service>vprn>sub-if>grp-if

Description This command enables remote proxy ARP on the interface.

Remote proxy ARP is similar to proxy ARP. It allows the router to answer an ARP request on an interface for a subnet that is not provisioned on that interface. This allows the router to forward to the other subnet on behalf of the requester. To distinguish remote proxy ARP from local proxy ARP, local proxy ARP performs a similar function but only when the requested IP is on the receiving interface.

The **no** form of this command reverts to the default.

3.6.2.1.7 Show Commands

Note: The command output in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

dhcp

Syntax **dhcp**

Context show>router
show>service>id

Description This command enables the context do display DHCP related information.

lease-state

Syntax **lease-state** [**wholesaler** *service-id*] [**sap** *sap-id*|**sdp** *sdp-id:vc-id*| **interface** *interface-name*|**ip-address** *ip-address*[*/mask*] | **chaddr** *ieee-address* | **mac** *ieee-address* | {[**port** *port-id*] [**no-inter-dest-id** | **inter-dest-id** *inter-dest-id*]}] [**session** {**none** | **ipoe**}] [**detail**]

Context show>service>id>dhcp

Description This command displays DHCP lease state information. The **wholesaler** *service-id* parameter is applicable only in the VPRN context.

Note: The **wholesaler** *service-id* parameter is applicable only in the VPRN context (VPRN is supported by the 7750 SR only).

Parameters *service-id* — Specifies the service ID of the wholesaler. When specified in this context, SAP, SDP, interface, IP address and MAC parameters are ignored. This parameter applies only to the 7750 SR.

Values service-id: 1 to 214748364
svc-name: A string, up to 64 characters

sap-id — Specifies the physical port identifier portion of the SAP definition

sdp-id — Specifies the SDP identifier

Values 1 to 17407

vc-id — Specifies the virtual circuit identifier. This value is used to validate the VC ID portion of each mesh SDP binding defined in the service. The default value of this object is equal to the service ID.

Values 1 — 4294967295

interface-name — Shows the interface name, up to 32 characters

ip-address[/mask] — Shows information for the specified IP address and mask

port-id — Specifies the DHCP lease state local specifies that the DHCP lease state is learned by either a SAP or SDP. When the value is SAP, the value indicates the SAP for which this entry contains information.

chaddr — Specifies the MA address of the DHCP lease state

interface-name — Shows information for the specified IP interface

detail — Displays detailed lease state information

inter-dest-id — Indicates the intermediate destination identifier received from either the DHCP or the RADIUS server or the local user database

session — Shows DHCPv4 lease states for hosts that are associated with an IPoE session or for hosts that are not associated with an IPoE session

Values none, ipoe

detail — Displays detailed information

Output The following output is an example of lease state information.

Sample Output

```
*A:ALA-48>config# show service id 101 dhcp lease-state
===============================================================================
DHCP lease state table, service 101
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LifeTime   Origin   Stdby
-------------------------------------------------------------------------------
10.1.1.52      00:00:1f:bd:00:bb lag-1:101           00h02m56s  DHCP-R
10.3.2.62      00:00:1f:bd:00:c6 lag-1:105           00h02m59s  RADIUS
-------------------------------------------------------------------------------
Number of lease states : 2
===============================================================================
*A:ALA-48>config#


*A:ALA-48>config# show service id 105 dhcp lease-state  wholesaler 101
===============================================================================
DHCP lease state table, service 105
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                      LifeTime   Origin   Stdby
-------------------------------------------------------------------------------
Wholesaler 101 Leases
-------------------------------------------------------------------------------
10.3.2.62      00:00:1f:bd:00:c6 lag-1:105           00h00m39s  RADIUS
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
*A:ALA-48>config#
```

Table 40 describes the output fields for DHCP summary.

Table 40 **DHCP Summary Field Descriptions**

Label	Description
Interface Name	The name of the router interface
ARP Populate	The ARP populate is enabled or disabled
Used/Provided	The number of used and provided DHCP leases
Info Option	Option 82 processing is enabled or disabled on the interface
Admin State	The administrative state

dhcp

Syntax **dhcp**

Context tools>perform>router

Description This command enables the context to configure tools DHCP parameters.

dhcp6

Syntax **dhcp6**

Context tools>perform>router

Description This command enables the context to configure tools DHCP6 parameters.

local-dhcp-server

Syntax **local-dhcp-server** *server-name*

Context show>router>dhcp
show>router>dhcp6
tools>perform>router>dhcp>server>pool

Description This command displays local DHCP or DHCP6 server information.

Parameters *server-name* — Specifies information about the local DHCP server

associations

Syntax **associations**

Context show>router>dhcp>server
show>router>dhcp
show>router>dhcp6

Description This command displays the interfaces associated with this DHCP or DHCP6 server.

Output The following output is an example of information associated with this DHCP or DHCP6 server.

Sample Output

```
*A:SUB-Dut-A# show router dhcp local-dhcp-server dhcpS1 associations
===============================================================================
DHCP server s1  router 3
===============================================================================
Associations                    Admin
-------------------------------------------------------------------------------
tosim5                          Up
===============================================================================
*A:SUB-Dut-A#
```

Table 41 describes DHCP associations fields.

Table 41 **DHCP6 Summary Field Descriptions**

Label	Description
Associations	The entity associated with the local DHCP server
Admin	The administrative state

declined-addresses

Syntax **declined-addresses** *ip-address*[/*mask*] [**detail**]
declined-addresses pool *pool-name*

Context show>router>dhcp>server

Description This command display information about declined addresses.

Parameters *pool-name* — Specifies a DHCP pool name on the router

ip-address — Specifies the IP address of the DNS server. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

detail — Displays detailed information

Output The following output is an example of declined address information

Sample Output

```
*A:ALA-48>show>router>dhcp>server# declined-addresses pool test
===============================================================================
Declined addresses for server test Base
===============================================================================
 Pool                              Subnet              IP Address
PPPoe User Name                  Time                MAC Address       Type
Option 82 Circuit ID
-------------------------------------------------------------------------------
No Matching Entries
===============================================================================
*A:ALA-48>show>router>dhcp>server#
```

Table 42 describes declined address pool output fields.

Table 42 **Declined Address Pool Field Descriptions**

Field	Description
Pool	The name of the IP address pool
Subnet	The subnet address
IP Address	The declined IP address
Time	The creation time of this entry
MAC Address	The declined MAC address
Type	The type of declined address
PPPoe User Name	The name of the PPPoE user name
Option 82 Circuit ID	The declined circuit ID from the Option 82

failover-pool-stats

Syntax **failover-pool-stats** [*pool-name*]

Context show>router>dhcp>local-dhcp-server
show>router>dhcp6>local-dhcp-server

Description This command displays failover pool statistics.

Parameters *pool-name* — Specifies the pool name which is defined in the system

Output The following output is an example of failover pool stats information

Sample Output

```
*A:vsim-2# show router 500 dhcp6 local-dhcp-server "d6" failover-pool-stats
===============================================================
Failover config for pool         v6-1
===============================================================
Failover Admin State             inService
Failover Oper State              preNormal
Failover Persist Key             0xFFFFFFFF
Administrative MCLT              0h10m0s
Operational MCLT                 0h10m0s
Startup Wait Time                0h2m0s
Partner Down Delay               23h59m59s
  Ignore MCLT                    disabled
Failover statistics for pool     v6-1
---------------------------------------------------------------
Dropped Invalid Packets          0
Failover Shutdown                0
Lease Already Expired            0
Maximum Lease Count Reached      0
Subnet Not Found                 0
Range Not Found                  0
Host Conflict                    0
Address Conflict                 0
Peer Conflict                    0
Persistence Congestion           0
No Lease Hold Time Configured    0
Lease Not Found                  0
---------------------------------------------------------------
Number of pools found            1
===============================================================
```

Table 43 describes failover pool statistics output fields.

Table 43 **Failover Pool Statistics Field Descriptions**

Field	Description
Failover Admin State	Identifiers the failover state of the DHCP server instance inService — The failover admin state is in service outOfService — The failover admin state is out of service
Failover Oper State	The operational state of a DHCP server instance
Failover Persist Key	The maximum amount of time that one server can extend a lease for a client's binding beyond the time known by the partner server
Administrative MCLT	The administrative Maximum Client Lead Time (MCLT)
Operational MCLT	Indicates the operational MCLT

Table 43 **Failover Pool Statistics Field Descriptions (Continued)**

Field	Description
Startup Wait Time	The startup wait time. The startup wait time is the time that one IP address pool attempts to contact the partner IP address pool. During this time, the IP address pool is unresponsive to DHCP client requests.
Partner Down Delay	The minimum safe-time after the beginning of COMMUNICATIONS-INTERRUPTED state. After the expiry of this time, the server enters the PARTNER-DOWN state.
Ignore McLT	The ignore McLT status. If, after the transition COMMUNICATIONS-INTERRUPTED-to-PARTNER-DOWN state, the DHCP server instance ignores the safety period with a duration of Maximum Client Lead Time; a ‘true’ value has the effect that the DHCP server starts offering IP addresses from the partner's scope immediately after this transition, without waiting for existing leases allocated by the partner and not known by this system to time out. A ‘true’ value increases the risk that duplicate addresses are offered; if the transition to PARTNER-DOWN state is likely to be caused by a failure of the partner system rather than a communications problem, this risk is reduced.
Failover statistics for pool	The failover statistics for each pool
Dropped Invalid Packets	The number of BNDUPD packets that were dropped because the packet was malformed
Failover Shutdown	The number of BNDUPD packets that were dropped because the failover state if the DHCP server instance is shut down
Lease Already Expired	The number of BNDUPD packets that were dropped because the corresponding lease has expired
Maximum Lease Count Reached	The number of BNDUPD packets that were dropped because the maximum number of leases were reached
Subnet Not Found	The number of BNDUPD packets that were dropped because a valid subnet could not be found for the lease
Range Not Found	The number of BNDUPD packets that were dropped because a valid include range could not be found for the lease

Table 43 **Failover Pool Statistics Field Descriptions (Continued)**

Field	Description
Host Conflict	The number of BNDUPD packets that were dropped because this DHCP server instance has already leased this address to another host
Address Conflict	The number of BNDUPD packets that were dropped because this DHCP server instance has already leased another address to this host
Peer Conflict	The number of BNDUPD packets that were dropped because the failover peer has leased an address within a subnet range of which the failover control is set to **local** on this local DHCP server instance
Persist Congestion	The number of BNDUPD packets that were dropped because of persistence congestion on this DHCP server instance
No Lease Hold Time Configured	The number of BNDUPD packets that were dropped because the lesae hold time is zero on this DHCP server instance
Lease Not Found	The number of Binding Database Update (BNDUPD) remove packets were dropped because the corresponding lease could not be found.
Number of pools found	The total number of pools found

failover-server-stats

Syntax **failover-server-stats**

Context show>router>dhcp>local-dhcp-server
show>router>dhcp6>local-dhcp-server

Description This command displays failover server statistics.

Output The following output is an example of failover server stats information.

Sample Output

```
show>router>dhcp# local-dhcp-server "y" failover-server-stats
===============================================================
Failover config for server       y
===============================================================
Failover Admin State             outOfService
Failover Oper State              shutdown
Failover Persist Key             0xFFFFFFFF
Administrative MCLT              0h10m0s
```

```
Operational MCLT                0h10m0s
Startup Wait Time               0h2m0s
Partner Down Delay              23h59m59s
  Ignore MCLT                   disabled

Failover statistics for server   y
-------------------------------------------------------------------
Dropped Invalid Packets          0
Failover Shutdown                0
Lease Already Expired            0
Maximum Lease Count Reached      0
Subnet Not Found                 0
Range Not Found                  0
Host Conflict                    0
Address Conflict                 0
Peer Conflict                    0
Persistence Congestion           0
No Lease Hold Time Configured    0
Invalid Prefix Length            0
Lease Not Found                  0
===================================================================
```

Table 44 describes failover server stats fields.

Table 44 **Failover Server Statistics Field Descriptions**

Field	Description
Failover Admin State	Identifiers the failover state of the DHCP server instance inService — The failover admin state is in service outOfService — The failover admin state is out of service
Failover Oper State	The operational state of a DHCP server instance
Failover Persist Key	The maximum amount of time that one server can extend a lease for a client's binding beyond the time known by the partner server
Administrative MCLT	The administrative Maximum Client Lead Time (MCLT)
Operational MCLT	Indicates the operational MCLT
Startup Wait Time	The startup wait time. The startup wait time is the time that one IP address pool attempts to contact the partner IP address pool. During this time, the IP address pool is unresponsive to DHCP client requests.
Partner Down Delay	The minimum safe-time after the beginning of COMMUNICATIONS-INTERRUPTED state. After the expiry of this time, the server enters the PARTNER-DOWN state.

Table 44 **Failover Server Statistics Field Descriptions (Continued)**

Field	Description
Ignore McLT	The ignore McLT status. If, after the transition COMMUNICATIONS-INTERRUPTED-to-PARTNER-DOWN state, the DHCP server instance ignores the safety period with a duration of Maximum Client Lead Time; a 'true' value has the effect that the DHCP server starts offering IP addresses from the partner's scope immediately after this transition, without waiting for existing leases allocated by the partner and not known by this system to time out. A 'true' value increases the risk that duplicate addresses are offered; if the transition to PARTNER-DOWN state is likely to be caused by a failure of the partner system rather than a communications problem, this risk is reduced.
Failover statistics for server	The failover statistics for each pool
Dropped Invalid Packets	The number of BNDUPD packets that were dropped because the packet was malformed
Failover Shutdown	The number of BNDUPD packets that were dropped because the failover state if the DHCP server instance is shut down
Lease Already Expired	The number of BNDUPD packets that were dropped because the corresponding lease has expired
Maximum Lease Count Reached	The number of BNDUPD packets that were dropped because the maximum number of leases were reached
Subnet Not Found	The number of BNDUPD packets that were dropped because a valid subnet could not be found for the lease
Range Not Found	The number of BNDUPD packets that were dropped because a valid include range could not be found for the lease.
Host Conflict	The number of BNDUPD packets that were dropped because this DHCP server instance has already leased this address to another host
Address Conflict	The number of BNDUPD packets that were dropped because this DHCP server instance has already leased another address to this host
Peer Conflict	The number of BNDUPD packets that were dropped because the failover peer has leased an address within a subnet range of which the failover control is set to **local** on this local DHCP server instance

Table 44 **Failover Server Statistics Field Descriptions (Continued)**

Field	Description
Persist Congestion	The number of BNDUPD packets that were dropped because of persistence congestion on this DHCP server instance
No Lease Hold Time Configured	The number of BNDUPD packets that were dropped because the lease hold time is zero on this DHCP server instance
Lease Not Found	The number of Binding Database Update (BNDUPD) remove packets were dropped because the corresponding lease could not be found.

free-addresses

Syntax **free-addresses** *ip-address*[/*mask*]
free-addresses summary [**subnet** *ip-address*[/*mask*]
free-addresses pool *pool-name*

Context show>router>dhcp>local-dhcp-server

Description This command displays the free addresses in a subnet.

Parameters *pool-name* — Specifies a DHCP pool name on the router

subnet — Specifies a subnet of IP addresses that are served from the pool.

summary — Displays summary output of the free addresses

Output The following output is an example of free address information

Sample Output

```
*A:ALA-48>show>router>dhcp>local-dhcp-server# free-
addresses pool test subnet  10.0.0.0/24
===============================================================================
Free addresses in subnet 10.0.0.0/24
===============================================================================
IP Address
-------------------------------------------------------------------------------
No. of free addresses: 0
===============================================================================
*A:ALA-48>show>router>dhcp>local-dhcp-server#
```

Table 45 describes free addresses pool command output fields.

Table 45 **Free Addresses Field Descriptions**

Field	Description
IP Address	The free IP addresses in the subnet
No of free addresses	The total number of free addresses in the subnet

leases

Syntax **leases** [**detail**]
leases *ip-address*[/*mask*] **address-from-user-db** [**detail**]
leases *ip-address*[/*mask*] **dhcp-host** *dhcp-host-name* [**detail**]
leases *ip-address*[/*mask*] **ppp-host** *ppp-host-name* [**detail**]
leases *ip-address*[/*mask*] [**detail**]

Context show>router>dhcp>local-dhcp-server

Description This command displays the DHCP leases.

Parameters *ip-address* — Specifies the base IP address of the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

mask — The subnet mask in dotted decimal notation

Values 0 to 32

address-from-user-db — Displays only leases that have IP addresses from the local-user-db

dhcp-host-name — Shows the leases that match a certain DHCP host from the local-user-db

ppp-host-name — Displays the leases that match a certain PPPoE host from the local-user-db

detail — Displays detailed information of all leases that fall into the indicated subnet.

The command with no parameters shows all leases from the local-user-db.

Output The following output is an example of lease information

Sample Output

```
*A:ALA-48>show>router>dhcp>local-dhcp-server# leases ip-address 10.0.0.4
===============================================================================
Leases for DHCP server test router Base
===============================================================================
IP Address      Lease State       Mac Address       Remaining Clnt
  PPPoE user name/Opt82 Circuit Id                  LifeTime  Type
-------------------------------------------------------------------------------
No leases found
```

```
*A:ALA-48>show>router>dhcp>local-dhcp-server#
```

Table 46 describes leases output fields.

Table 46 **Leases Field Descriptions**

Field	Description
IP Address	The IP address for the specified lease
Lease State	The lease state of the IP address
Mac Address	The MAC address of the lease
Remaining Lifetime	The remaining lifetime of the lease
Clnt Type	The address type
PPPoE user name	The PPPoE user name
Opt82	The option number that the DHCP server uses to send the identification strings to the PPPoE client
Circuit ID	The circuit ID from Option 82

pool-ext-stats

Syntax **pool-ext-stats** [*pool-name*]

Context show>router>dhcp>server

Description This command displays extended statistics per DHCPv4 pool in local DHCPv4 server.

The following statistics are included in output:

- The number of stable leases in the pool
- The number of provisioned address in the pool
- The number of used address in the pool
- The number of free address in the pool
- The percentage of used address
- The percentage of free address

For each statistic (except for provisioned addresses), there is current value and peak value, peak value is the highest value since pool creation or last reset via the **clear router** *rt-id* **dhcp local-dhcp-server** *svr-name* **pool-ext-stats** command.

Parameters *pool-name* — Specifies the name of DHCPv4 local server pool

Output The following output is an example of pool extended statistics information.

Sample Output

```
show router 500 dhcp local-dhcp-server "d4" pool-ext-stats "pool-1"
===============================================================================
Extended pool statistics for server "d4"
===============================================================================
                          Current         Peak            TimeStamp
-------------------------------------------------------------------------------
Pool                      pool-1
Local:
  Stable Leases           0               0               01/07/2013 19:07:11
  Provisioned Addresses   101
  Used Addresses          0               0               01/07/2013 19:07:11
  Free Addresses          101             101             01/07/2013 19:07:11
  Used Pct                0               0               01/07/2013 19:07:11
  Free Pct                100             100             01/07/2013 19:07:11
Last Reset Time                                           01/07/2013 19:07:11
-------------------------------------------------------------------------------
Number of entries         1
===============================================================================
```

Table 47 describes extended pool statistics output fields.

Table 47 **Extended Pool Statistics Field Descriptions**

Field	Description
Current	The current pool statistics
Peak	The peak value since the last reset
Timestamp	The date and time of the last reset
Pool	The pool name
Stable Leases	The number of stable leases in the pool
Provisioned Addresses	The number of provisioned addresses in the pool
Used Addressed	The number of used addresses in the pool
Free Addresses	The number of free addresses in the pool
Used Pct	The percentage of subnets in use in the pool
Free Pct	The percentage of unused subnets in the pool
Last Reset Time	The date and time of the last reset in the pool
Number of entries	The total number of entries

server-stats

Syntax **server-stats**

Context show>router>dhcp>server
show>router>dhcp6>server

Description This command displays DHCP or DHCP6 server statistics.

Output The following output is an example of server statics information

Sample Output

```
*A:SUB-Dut-A# show router dhcp local-dhcp-server dhcpS1 server-stats
===============================================================================
Statistics for DHCP Server dhcpS1 router Base
===============================================================================
Rx Discover Packets           : 0
Rx Request Packets            : 0
Rx Release Packets            : 0
Rx Decline Packets            : 0
Rx Inform Packets             : 0

Tx Offer Packets              : 0
Tx Ack Packets                : 0
Tx Nak Packets                : 0
Tx Forcerenew Packets         : 0

Client Ignored Offers         : 0
Leases Timed Out              : 0

Dropped Bad Packet            : 0
Dropped Invalid Type          : 0
Dropped No User Database      : 0
Dropped Unknown Host          : 0
Dropped User Not Allowed      : 0
Dropped Lease Not Ready       : 0
Dropped Lease Not Found       : 0
Dropped Not Serving Pool      : 0
Dropped Invalid User          : 0
Dropped Overload              : 0
Dropped Persistence Overload  : 0
Dropped Generic Error         : 0
Dropped Destined To Other     : 0
Dropped Address Unavailable   : 0
Dropped Max Leases Reached    : 0
Dropped Server Shutdown       : 0
Dropped No Subnet For Fixed IP: 0

===============================================================================
*A:SUB-Dut-A#
```

Table 48 describes server stats show command output fields.

Table 48 **Server Stats Field Descriptions**

Field	Description
RX Discover Packets	The number of discover packets received by the DHCP server instance
Rx Request Packets	The number of requst packets received by the DHCP server instance.
Rx Release Packets	The number of released packets received by the DHCP server instance
Rx Decline Packets	The number of declined packets received by the DHCP server instance
Rx Inform Packets	The number of inform packets received by the DHCP server instance
Tx Offer Packets	The number of offer packets sent by the DHCP server instance
Tx Ack Packets	The number of DHCP ACK packets sent by the DHCP server instance
Tx Nak Packets	The number of DHCP NAK packets sent by the DHCP server instance
Tx Forcerenew Packets	The number of DHCP force renew packets sent by the server instance
Client Ignored Offers	The number of DHCP offer packets sent by the DHCP server instance that were ignored by clients
Leases Timed Out	The number of leases timed out
Dropped Bad Packet	The number of dropped of DHCP packets received that were corrupt
Dropped Invalid Type	The number of DHCP packets received that had an invalid message type
Dropped No User Database	The number of DHCP packets dropped because the value of this server instance is not equal to the default value and a local user database with that name could not be found
Dropped Unknown Host	The number of DHCP packets dropped from hosts which were not found in the user database
Dropped User Not Allowed	The number of DHCP packets dropped from hosts which are found in the user database, but which have no address or pool specified

Table 48 **Server Stats Field Descriptions (Continued)**

Field	Description
Dropped Lease Not Ready	The number of DHCP packets dropped by the server instance before the lease database was ready
Dropped Lease Not Found	The number of DHCP packets dropped by the server instance because no (valid) lease was found
Dropped Not Serving Pool	The number of DHCP packets dropped by the server instance because there were no more free addresses in the pool
Dropped Invalid User	The number of DHCP packets dropped by the server instance because the MAC address of the sender or the option 82 didn't match the host lease state
Dropped Overload	The number of DHCP packets dropped by the server instance because they were received in excess of what the server instance can handle
Dropped Persistence Overload	The number of DHCP packets dropped by the server instance because they were received in excess of what the DHCP persistence system can handle. If this occurs, only releases and declines are still processed.
Dropped Generic Error	The number of DHCP packets dropped by the server instance because of a generic error
Dropped Destined To Other	The number of DHCP requests dropped by the server instance because the (broadcast) request was not destined to this server
Dropped Address Unavailable	The number of DHCP requests dropped by the server instance because the requested address is not available
Dropped Max Leases Reached	The number of DHCP packets dropped by the server instance because the maximum number of leases was reached
Dropped Server Shutdown	The number of DHCP packets dropped by the server instance during server instance shutdown
Dropped No Subnet for Fixed IP	The number of DHCP packets dropped by the server instance for user-db hosts with a fixed address because the subnet to which the address belongs is not configured

sticky-leases

Syntax **sticky-leases** [*hostname*]

Context show>router>dhcp>server

Description This command shows either all sticky leases or a single specific sticky lease created under the specified DHCP server.

Parameters *hostname* — Specifies the name of the host for which an entry is created

Output The following output is an example of sticky lease information

Sample Output

```
Node# show router dhcp local-dhcp-server "dhcp_server" sticky-leases
===============================================================================
DHCP Server Sticky Leases
===============================================================================
Hostname                          IP address
-------------------------------------------------------------------------------
test_lease                        10.0.0.1
-------------------------------------------------------------------------------
No. of Matching Entries: 1
===============================================================================
```

Table 49 describes the sticky leases field descriptions.

Table 49 **Sticky Leases Field Descriptions**

Field	Description
Hostname	The name of the sticky host
IP address	The IP address of the sticky host
No. of Matching Entries	The total numbers of entries matching the command criteria

subnet-ext-stats

Syntax **subnet-ext-stats** *ip-address*[/*mask*]
subnet-ext-stats pool *pool-name*

Context show>router>dhcp>servers

Description This command displays extended statistics per DHCPv4 subnet in local DHCPv4 server.

The following statistics are included in output:

- The number of stable leases in the subnet
- The number of provisioned address in the subnet
- The number of used address in the subnet
- The number of free address in the subnet

- The percentage of used address
- The percentage of free address

For each statistic (except for Provisioned Addresses), there is current value and peak value, peak value is the highest value since subnet creation or last reset via the **clear router** *rt-id* **dhcp local-dhcp-server** *svr-name* **subnet-ext-stats** command.

When parameter pool is used, the statistics of each subnet in the pool will be displayed.

Parameters *ip-address[/mask]* — Specifies the subnet and mask

pool-name — Specifies the name of local DHCPv4 server pool

Output The following output is an example of subnet extended statistics information.

Sample Output

```
show router 500 dhcp local-dhcp-server "d4" subnet-ext-stats 239.10.10.0/24
===============================================================================
Extended statistics for subnet 239.10.10.0/24
===============================================================================
                            Current         Peak            TimeStamp
-------------------------------------------------------------------------------
Local:
  Stable Leases             1               1               01/07/2013 19:38:36
  Provisioned Addresses     101
  Used Addresses            1               1               01/07/2013 19:38:36
  Free Addresses            100             100             01/07/2013 19:38:36
  Used Pct                  1               1               01/07/2013 19:38:36
  Free Pct                  99              99              01/07/2013 19:38:36
Last Reset Time                                             01/07/2013 19:07:11
-------------------------------------------------------------------------------
Number of entries           1
===============================================================================
```

Table 50 describes extended subnet statistics output field descriptions.

Table 50 **Extended Subnet Statistics Field Descriptions**

Field	Descriptions
Current	The current statistics for the subnet
Peak	The peak statistics for the subnet
TimeStamp	The timestamp for the last reset
Stable Leases	The number of stable leases
Provision Addresses	The number of provisioned addresses in this subnet
Used Addresses	The number of used addresses in this subnet
Free Address	The number of free addresses in this subnet

Table 50 **Extended Subnet Statistics Field Descriptions (Continued)**

Field	Descriptions
Used Pct	The percentage of addresses in this subnet in use
Free Pct	The percentage of addresses in this subnet currently unused
Last Reset Time	The time of the last reset
Number of entries	The total number of entries

subnet-stats

Syntax **subnet-stats** *ip-address*[/*mask*]
subnet-stats pool *pool-name*

Context show>router>dhcp>server

Description This command displays subnet statistics.

Parameters *ip-address[/mask]* — Specifies the subnet and mask

pool-name — Specifies the name of local DHCPv4 server pool

Output The following output is an example of subnet statistics information.

Sample Output

```
*A:SUB-Dut-A# show router dhcp local-dhcp-server dhcpS2 subnet-stats pool POOL2
===============================================================================
Statistics for pool POOL2
===============================================================================
Subnet              Free          Offered      Stable
                      FRPending     RemPending   Declined
-------------------------------------------------------------------------------
10.0.0.0/8          16384         0            0
                     0             0            0
-------------------------------------------------------------------------------
No. of entries: 1
===============================================================================
*A:SUB-Dut-A#
```

Table 51 describe the fields for subnet stats pool output.

Table 51 **Subnet Stats Pool Field Descriptions**

Field	Description
Subnet	The subnet address

Table 51 **Subnet Stats Pool Field Descriptions (Continued)**

Field	Description
Free	The number of addresses in this subnet that are free
Offered	The number of leases in this subnet that are in state **offered**
Stable	The number of leases in this subnet that are in state **stable**
FRPending	The number of leases in this subnet that are in state **forceRenewPendin**g
RemPending	The number of leases in this subnet that are in state **removePending**
Declined	The number of addresses in this subnet that are declined
No. of entries	The total number of entries

summary

Syntax **summary**

Context show>router>dhcp>server

Description This command displays DHCP server summary information.

Output The following output is an example of DHCP summary information.

Sample Output

```
*A:SUB-Dut-A# show router dhcp local-dhcp-server dhcpS2 summary
===============================================================================
DHCP server dhcpS2  router Base
===============================================================================
dhcpS2-POOL2
Admin State            : inService
Persistency State      : ok
User Data Base         : N/A
Use gateway IP address : disabled
Send force-renewals    : disabled
-------------------------------------------------------------------------------
Pool name : POOL2
-------------------------------------------------------------------------------
Subnet             Free       Stable     Declined   Offered    Remove-pending
-------------------------------------------------------------------------------
10.0.0.0/8         16384      0          0          0          0

Totals for pool    16384      0          0          0          0
-------------------------------------------------------------------------------
Totals for server  16384      0          0          0          0
-------------------------------------------------------------------------------
```

```
Associations                    Admin
-------------------------------------------------------------------------------
No associations found
===============================================================================
*A:SUB-Dut-A#


*A:vsim-2# show router 500 dhcp local-dhcp-server "d4" summary
===============================================================================
DHCP server d4  router 500
===============================================================================
Admin State            : inService
Operational State      : inService
Persistency State      : shutdown
User Data Base         : N/A
Use gateway IP address : enabled (scope subnet)
Use pool from client   : disabled
Send force-renewals    : disabled
Creation Origin        : manual
Lease Hold Time        : 0h0m0s
Lease Hold Time For    : N/A
User-ident             : mac-circuit-id
Failover Admin State   : outOfService
Failover Oper State    : shutdown
Failover Persist Key   : N/A
Administrative MCLT    : 0h10m0s
Operational MCLT       : 0h10m0s
Startup wait time      : 0h2m0s
Partner down delay     : 23h59m59s
  Ignore MCLT          : disabled
-------------------------------------------------------------------------------
Pool name : v4-1
-------------------------------------------------------------------------------
Failover Admin State   : inService
Failover Oper State    : normal
Failover Persist Key   : N/A
Administrative MCLT    : 0h10m0s
Operational MCLT       : 0h10m0s
Startup wait time      : 0h2m0s
Partner down delay     : 23h59m59s
  Ignore MCLT          : disabled
-------------------------------------------------------------------------------
Subnet                 Free     %    Stable   Declined Offered  Rem-pend Drain
-------------------------------------------------------------------------------
10.20.20.0/24      (L) 10       90%  1        0        0        0        N
                   (R) N/A           0        N/A      N/A      N/A      N
Totals for pool        10       90%  1        0        0        0
-------------------------------------------------------------------------------
Totals for server      10       90%  1        0        0        0
-------------------------------------------------------------------------------
Interface associations
Interface                       Admin
-------------------------------------------------------------------------------
l1                              Up
-------------------------------------------------------------------------------
Local Address Assignment associations
Group interface                 Admin
-------------------------------------------------------------------------------
===============================================================================
```

```
*A:vsim-2#
```

servers

Syntax **servers**
servers all

Context show>router>dhcp

Description This command lists the local DHCP servers.

Parameters **all** — Displays a complete list of local DHCP servers

Output The following output is an example of local DHCP server information.

Sample Output

```
*A:ALA-49>show>router>dhcp# servers
===============================================================
Overview of DHCP Servers
===============================================================
Active Leases:      0
Maximum Leases:     159744

Router              Server                          Admin State
---------------------------------------------------------------
Router: Base        base_router_dhcp_server         outOfService
Service: 3          s1                              inService
===============================================================
*A:ALA-49>show>router>dhcp#
```

statistics

Syntax **statistics** [**interface** *ip-int-name*]]

Context show>router>dhcp

Description This command displays statistics for DHCP relay and DHCP snooping.

If no IP address or interface name is specified, then all configured interfaces are displayed.

If an IP address or interface name is specified, then only data regarding the specified interface is displayed.

Parameters *ip-int-name* | *ip-address* — Displays statistics for the specified IP interface

Output The following is an example of DHCP statistics information.

Sample Output

```
*A:PE-5# show router dhcp statistics
===============================================================
DHCP Global Statistics (Router: Base)
===============================================================
Rx Packets                          : 0
Tx Packets                          : 0
Rx Malformed Packets                : 0
Rx Untrusted Packets                : 0
Client Packets Discarded            : 0
Client Packets Relayed              : 0
Client Packets Snooped              : 0
Client Packets Proxied (RADIUS)     : 0
Client Packets Proxied (Diameter)   : 0
Client Packets Proxied (User-Db)    : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Discarded            : 0
Server Packets Relayed              : 0
Server Packets Snooped              : 0
DHCP RELEASEs Spoofed               : 0
DHCP FORCERENEWs Spoofed            : 0
Client packets streamed             : 0
===============================================================
```

Table 52 describes the output fields for DHCP statistics.

Table 52 **DHCP Statistics Fields**

Label	Description
Rx Packets	The number of packets received from the DHCP clients
Tx Packets	The number of packets transmitted to the DHCP clients
Rx Malformed Packets	The number of malformed packets received from the DHCP clients
Rx Untrusted Packets	The number of untrusted packets received from the DHCP clients
Client Packets Discarded	The number of packets received from the DHCP clients that were discarded
Client Packets Relayed	The number of packets received from the DHCP clients that were forwarded
Client Packets Snooped	The number of packets received from the DHCP clients that were snooped
Server Packets Discarded	The number of packets received from the DHCP server that were discarded
Server Packets Relayed	The number of packets received from the DHCP server that were forwarded
Server Packets Snooped	The number of packets received from the DHCP server that were snooped

statistics

Syntax **statistics**

Context show>router>dhcp6

Description This command displays statistics for DHCP relay and DHCP snooping.

If no IP address or interface name is specified, then all configured interfaces are displayed.

If an IP address or interface name is specified, then only data regarding the specified interface is displayed.

Output The following output is an example of DHCP6 statistics.

Sample Output

```
*A:PE-5# show router dhcp6 statistics
===============================================================================
DHCP6 statistics (Router: Base)
===============================================================================
Msg-type                      Rx             Tx             Dropped
-------------------------------------------------------------------------------
1 SOLICIT                     0              0              0
2 ADVERTISE                   0              0              0
3 REQUEST                     0              0              0
4 CONFIRM                     0              0              0
5 RENEW                       0              0              0
6 REBIND                      0              0              0
7 REPLY                       0              0              0
8 RELEASE                     0              0              0
9 DECLINE                     0              0              0
10 RECONFIGURE                0              0              0
11 INFO_REQUEST               0              0              0
12 RELAY_FORW                 0              0              0
13 RELAY_REPLY                0              0              0
14 LEASEQUERY                 0              0              0
15 LEASEQUERY_REPLY           0              0              0
-------------------------------------------------------------------------------
Dhcp6 Drop Reason Counters :
-------------------------------------------------------------------------------
1 Dhcp6 oper state is not Up on src itf                          0
2 Dhcp6 oper state is not Up on dst itf                          0
3 Relay Reply Msg on Client Itf                                  0
4 Hop Count Limit reached                                        0
5 Missing Relay Msg option, or illegal msg type                  0
6 Unable to determine destination client Itf                     0
7 Out of Memory                                                  0
8 No global Pfx on Client Itf                                    0
9 Unable to determine src Ip Addr                                0
10 No route to server                                            0
11 Subscr. Mgmt. Update failed                                   0
12 Received Relay Forw Message                                   0
13 Packet too small to contain valid dhcp6 msg                   0
14 Server cannot respond to this message                         0
15 No Server Id option in msg from server                        0
```

```
16 Missing or illegal Client Id option in client msg             0
17 Server Id option in client msg                                0
18 Server DUID in client msg does not match our own              0
19 Client sent message to unicast while not allowed              0
20 Client sent message with illegal src Ip address               0
21 Client message type not supported in pfx delegation           0
22 Nbr of addrs or pfxs exceeds allowed max (128) in msg         0
23 Unable to resolve client's mac address                        0
24 The Client was assigned an illegal address                    0
25 Illegal msg encoding                                          0
26 Client message not supported                                  0
27 IA options in info request                                    0
28 No IA option in client msg                                    0
29 No addresses in confirm msg                                   0
30 No relay servers configured                                   0
31 Blocked by host lockout                                       0
32 No link address available                                     0
33 Dropped by Python                                             0
34 Invalid server                                                0
35 Packet dropped on SRRP backup interface                       0
36 DHCP transaction not found                                    0
37 Could not determine retail interface                          0
38 Packet dropped by DHCP filter                                 0
39 Packet dropped because authentication failed                  0
40 Invalid forced source address                                 0
===============================================================================
```

Table 53 describers DHCP6 statistics output information.

Table 53 **DHCP6 Statistics Field Descriptions**

Field	Description
Msg-Type	A list of message types that lists basic statistics about a particular DHCP server instance.
Rx	The number of messages received
Tx	The number of messages sent
Dropped	The number of messages dropped
Dhcp6 Drop Reason Counters	The number of instances for the applicable drop reason

statistics

Syntax **statistics** [[**sap** *sap-id*] | [**sdp** *sdp-id:vc-id*] | [**interface** *interface-name*]]

Context show>service>id>dhcp

Description This command displays DHCP relay statistics.

Parameters *ip-int-name* — Displays DHCP statistics on the specified interface

interface-name — Displays DHCP statistics for the specified interface name

sap-id — Displays DHCP statistics for the specified SAP

Output The following output is an example of DHCP statistics information.

Sample Output

```
*A:ALA-48# show service id 88 dhcp statistics interface SpokeTerm
===============================================================================
DHCP Statistics for interface SpokeTerm
===============================================================================
Rx Packets                          : 0
Tx Packets                          : 0
Rx Malformed Packets                : 0
Rx Untrusted Packets                : 0
Client Packets Discarded            : 0
Client Packets Relayed              : 0
Client Packets Snooped              : 0
Client Packets Proxied (RADIUS)     : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Discarded            : 0
Server Packets Relayed              : 0
Server Packets Snooped              : 0
DHCP RELEASEs Spoofed               : 0
DHCP FORCERENEWs Spoofed            : 0
===============================================================================
*A:ALA-48#
```

Table 54 describes DHCP statistics output fields.

Table 54 **DHCP Statistics Field Descriptions**

Field	Description
Rx Packets	The number of received packets
Tx Packets	The number of transmitted packets
Rx Malformed	The number of malformed packets received
Rx Untrusted Packets	The number of untrusted packets received
Client Packets Discarded	The number of client packets discarded
Client Packets Relayed	The number of client packets relayed
Client Packets Snooped	The number of client packets snooped
Client Packets Proxied (RADIUS)	The number of client packets proxied through RADIUS
Client Packets Proxied (Lease-Split)	The number of client packets proxied through a lease split

Table 54 **DHCP Statistics Field Descriptions (Continued)**

Field	Description
Server Packets Discarded	The number of discarded server packets
Server Packets Relayed	The number of relayed server packets
Server Packets Snooped	The number of snooped server packets
DHCP RELEASEs Spoofed	The number of spoofed DHCP releases
DHCP FORCERENEWs Spoofed	The number of spoofed FORCERENEW messages

summary

Syntax **summary**

Context show>router>dhcp
show>router>dhcp6

Description This command displays the status of the DHCP6 relay and DHCP snooping functions on each interface.

Output The following outputs display DHCP and DHCP6 summary information.

Sample Output

```
A:ALA-48>show>router>dhcp# summary
===============================================================================
Interface Name                   Arp      Used/     Info    Admin
                                 Populate Provided  Option  State
-------------------------------------------------------------------------------
ccaiesif                         No       0/0       Keep    Down
ccanet6                          No       0/0       Keep    Down
iesBundle                        No       0/0       Keep    Up
spokeSDP-test                    No       0/0       Keep    Down
test                             No       0/0       Keep    Up
test1                            No       0/0       Keep    Up
test2                            No       0/0       Keep    Up
testA                            No       0/0       Keep    Up
testB                            No       0/0       Keep    Up
testIES                          No       0/0       Keep    Up
to-web                           No       0/0       Keep    Up
-------------------------------------------------------------------------------
Interfaces: 11
===============================================================================
A:ALA-48>show>router>dhcp#


*A:vsim-2# show router 500 dhcp summary
===============================================================================
DHCP Summary (Service: 500)
```

```
===============================================================================
Interface Name                   Arp      Leases Per Interface/ Info    Admin
  SapId/Sdp                      Populate Per Sap Limit         Option  State
-------------------------------------------------------------------------------
g1                               No       1/1                   Keep    Up
  sap:1/1/7                               1/1
l1                               No       0/0                   Keep    Down
-------------------------------------------------------------------------------
Interfaces: 2
===============================================================================
*A:vsim-2#
```

Table 55 describes DHCP summary output field descriptions.

Table 55 **DHCP Summary Field Descriptions**

Label	Description
Interface Name	The name of the router interface
ARP Populate	ARP populate is enabled or disabled
Used/Provided	The number of used and provided DHCP leases
Info Option	Option 82 processing is enabled or disabled on the interface
Admin State	The administrative state

Sample Output

```
*A:vsim-2# show router 500 dhcp6 summary
===============================================================================
DHCP6 Summary (Service: 500)
===============================================================================
-------------------------------------------------------------------------------
Service Interfaces
-------------------------------------------------------------------------------
Interface Name                   Nbr      Used/Max Relay     Admin  Oper Relay
  SapId                          Resol.   Used/Max Server    Admin  Oper Server
-------------------------------------------------------------------------------
l1                               No          0/0             Down   Down
                                             0/8000          Down   Down
-------------------------------------------------------------------------------
Interfaces: 1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Routed CO Interfaces
-------------------------------------------------------------------------------
Interface Name                           Used               Admin (proxy/relay)
  SapId
-------------------------------------------------------------------------------
g1                                       1                  Down/Up
  sap:1/1/7                              1
-------------------------------------------------------------------------------
Interfaces: 1                            1
-------------------------------------------------------------------------------
```

```
===============================================================================
*A:vsim-2#

*A:vsim-2# show router 500 dhcp6 local-dhcp-server "d6" summary
===============================================================================
DHCP server d6  router 500
===============================================================================
Admin State            : inService
Operational State      : inService
Persistency State      : shutdown
Use Link Address       : enabled (scope subnet)
Use pool from client   : disabled
Creation Origin        : manual
Lease Hold Time        : 0h0m0s
Lease Hold Time For    : N/A
User-ident             : duid
Interface-id-mapping   : disabled
Ignore-rapid-commit    : disabled
Allow-lease-query      : disabled
Failover Admin State   : outOfService
Failover Oper State    : shutdown
Failover Persist Key   : N/A
Administrative MCLT    : 0h10m0s
Operational MCLT       : 0h10m0s
Startup wait time      : 0h2m0s
Partner down delay     : 23h59m59s
  Ignore MCLT          : disabled

-------------------------------------------------------------------------------
Pool name : v6-1
-------------------------------------------------------------------------------
Failover Admin State   : inService
Failover Oper State    : preNormal
Failover Persist Key   : N/A
Time Left              : 0h4m16s before state transition
Administrative MCLT    : 0h10m0s
Operational MCLT       : 0h10m0s
Startup wait time      : 0h2m0s
Partner down delay     : 23h59m59s
  Ignore MCLT          : disabled
-------------------------------------------------------------------------------
Prefix
                                    Stable   Declined Advert   Rem-pend Drain
-------------------------------------------------------------------------------
2001:dead:1::/56
                                (L) 1        0        0        0        N
2001:dead:2::/56
                                (R) 0        0        0        0        N
Totals for pool                     1        0        0        0
-------------------------------------------------------------------------------
Totals for server                   1        0        0        0
-------------------------------------------------------------------------------
Interface associations
Interface                       Admin
-------------------------------------------------------------------------------
l1                              Up
-------------------------------------------------------------------------------
Local Address Assignment associations
Group interface                 Admin
```

```
-------------------------------------------------------------------------------
===============================================================================
*A:vsim-2#
```

Table 56 describes DHCP6 summary output fields.

Table 56 **DHCP6 Summary Fields**

Label	Description
Interface Name	The name of the router interface
ARP Populate	ARP populate is enabled or disabled
Used/Provided	The number of used and provided DHCP leases
Info Option	Option 82 processing enabled or disabled on the interface
Admin State	The administrative state

dhcp6

Syntax **dhcp6**

Context show>router
show>service>id
show>system

Description This command enables the context to display DHCP6 related information.

interface-id-mapping

Syntax **interface-id-mapping**

Context show>router>dhcp6>local-dhcp-server

Description This command displays information about the DHCPv6 server that uses a unique /64 prefix per interface-ID options combination.

Output The following output is an example of interface ID mapping information.

Sample Output

```
show router 600 dhcp6 local-dhcp-server "d6" interface-id-mapping
===============================================================================
Interface-ID Mappings for DHCPv6 server d6
===============================================================================
Mapped Prefix      : 2001:AAAA::/64
Relay Interface ID : 1/1/10
LDRA Interface ID  : (Not Specified)
```

```
Active Leases      : 2001:AAAA::1 (stable)
===============================================================================
1 prefix found
===============================================================================
```

Table 57 describes interface ID mapping field descriptions.

Table 57 **Interface ID Mapping Field Descriptions**

Field	Description
Mapped Prefix	The mapped prefix
Relay Interface ID	The value of the Interface-ID Option assigned by the DHCPv6 Relay Agent
LDRA Interface ID	The value of the Interface-ID Option assigned by the Lightweight DHCPv6 Relay Agent
Active Leases	The number of active leases
No. of prefixes found	The total number of prefixes found.

leases

Syntax **leases** [*ipv6-address/prefix-length*] [*type*] [*state*] [**detail**]

Context show>router>dhcp6>local-dhcp-server

Description This command displays the DHCP6 leases.

Parameters *ipv6-address* — Specifies the base IP address of the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

mask — The subnet mask in dotted decimal notation

Values 0 to 32

type — Displays the lease type

Values pd, wan-host

state — Displays the state of the lease

Values advertised, remove-pending, held

detail — Displays detailed information of all leases that fall into the indicated subnet. The command with no parameters shows all leases from the local user database.

Output The following output is an example of lease information.

Sample Output

```
show router 600 dhcp6 local-dhcp-server "d6" leases
===============================================================================
Leases for DHCPv6 server d6
===============================================================================
IP Address/Prefix                          Lease State       Remaining   Fail
  Link-local Address                                         LifeTime    Ctrl
-------------------------------------------------------------------------------
2001:AAAA::1/128
  FE80::220:FCFF:FE1E:CD52                 stable            23h58m52s   local
-------------------------------------------------------------------------------
1 leases found
===============================================================================
```

Table 58 describes DHCP6 lease fields.

Table 58 **DHCP6 Lease Field Descriptions**

Field	Description
IP Address/Prefix	The IP address and prefix of the lease
Link-local Address	The link local address of the lease
Lease State	The lease state of the lease
Remaining Lifetime	The remaining lifetime of the lease
Fail Ctrl	The type of failover control
No. leases found	The total number of leases

pool-ext-stats

Syntax **pool-ext-stats** [*pool-name*]

Context show>router>dhcp6>server

Description This command displays extended statistics per DHCPv6 pool in local DHCPv6 server.

The following statistics are included in output:

- The number of stable leases in the pool
- The number of provisioned /64 address block in the pool
- The number of used /64 address block in the pool
- The number of free /64 address block in the pool
- The percentage of used address (with /64 address block)
- The percentage of free address (with /64 address block)

For each statistic (except for Provisioned Addresses), there is current value and peak value, peak value is the highest value since pool creation or last reset via command “clear router <rt-id> dhcp6 local-dhcp-server <svr-name> pool-ext-stats”.

Parameters *pool-name* — Specifies the name of DHCPv6 local server pool

Output The following output is an example of subnet statistics information.

Sample Output

```
show router 500 dhcp6 local-dhcp-server "d6" pool-ext-stats "pool-v6"
===============================================================================
Extended pool statistics for server "d6"
===============================================================================
                           Current         Peak            TimeStamp
-------------------------------------------------------------------------------
Pool                       pool-v6
Local:
  Stable Leases            0               0               01/07/2013 19:54:52
  Provisioned Blks         4
  Used Blks                0               0               01/07/2013 19:54:52
  Free Blks                4               4               01/07/2013 19:54:52
  Used Pct                 0               0               01/07/2013 19:54:52
  Free Pct                 100             100             01/07/2013 19:54:52
Last Reset Time                                            01/07/2013 19:54:52
-------------------------------------------------------------------------------
Number of entries          1
===============================================================================
```

Table 59 describes DHCP6 extended pool statistics output fields.

Table 59 **DHCP6 Pool Ext Stats Field Descriptions**

Field	Description
Current	The current pool statistics
Peak	The peak value since the last reset
TimeStamp	The date and time of the last reset
Pool	The pool name
Stable Leases	The number of stable leases
Provisioned Blks	The number of provisioned blocks in this pool
Free Blks	The number of free blocks in this pool
Used Pct	The percentage of extended statistics in use
Free Pct	The percentage of /64 blocks currently unused
Last Reset Time	The date and time of the last reset

Table 59 **DHCP6 Pool Ext Stats Field Descriptions (Continued)**

Field	Description
Number of entries	The total number of entries

pool-stats

Syntax **pool-stats** [*pool-name*]

Context show>router>dhcp6>server

Description This command displays pool statistics.

Parameters *pool-name* — Specifies the name of DHCPv6 local server pool

Output The following output is an example of pool statistics information.

Sample Output

```
*A:vsim-2# show router 500 dhcp6 local-dhcp-server "d6" pool-stats
===============================================================================
DHCPv6 server pool statistics
===============================================================================
Pool                          : v6-1
-------------------------------------------------------------------------------
Dropped Int no prefix WAN     : 0
Dropped Int no prefix SLAAC   : 0
===============================================================================
*A:vsim-2#
*A:vsim-2# show router 500 dhcp local-dhcp-server "d4" failover-
failover-pool-stats     failover-server-stats
*A:vsim-2# show router 500 dhcp local-dhcp-server "d4" failover-pool-stats
===============================================================
Failover config for pool         v4-1
===============================================================
Failover Admin State             inService
Failover Oper State              normal
Failover Persist Key             0xFFFFFFFF
Administrative MCLT              0h10m0s
Operational MCLT                 0h10m0s
Startup Wait Time                0h2m0s
Partner Down Delay               23h59m59s
  Ignore MCLT                    disabled
Failover statistics for pool     v4-1
---------------------------------------------------------------
Dropped Invalid Packets          0
Failover Shutdown                0
Lease Already Expired            0
Maximum Lease Count Reached      0
Subnet Not Found                 0
Range Not Found                  0
Host Conflict                    0
Address Conflict                 0
```

```
Peer Conflict                    0
Persistence Congestion           0
No Lease Hold Time Configured    0
Lease Not Found                  0
-------------------------------------------------------------------
Number of pools found            1
===================================================================
*A:vsim-2#
```

Table 60 describes pool statistics output field descriptions.

Table 60 **Failover Pool Statistics Field Descriptions**

Field	Description
Failover Admin State	Identifiers the failover state of the DHCP server instance inService — The failover admin state is in service outOfService — The failover admin state is out of service
Failover Oper State	The operational state of a DHCP server instance
Failover Persist Key	The maximum amount of time that one server can extend a lease for a client's binding beyond the time known by the partner server
Administrative MCLT	The administrative Maximum Client Lead Time (MCLT)
Operational MCLT	Indicates the operational MCLT
Startup Wait Time	The startup wait time. The startup wait time is the time that one IP address pool attempts to contact the partner IP address pool. During this time, the IP address pool is unresponsive to DHCP client requests.
Partner Down Delay	The minimum safe-time after the beginning of COMMUNICATIONS-INTERRUPTED state. After the expiry of this time, the server enters the PARTNER-DOWN state.

Table 60 **Failover Pool Statistics Field Descriptions (Continued)**

Field	Description
Ignore McLT	The ignore McLT status. If, after the transition COMMUNICATIONS-INTERRUPTED-to-PARTNER-DOWN state, the DHCP server instance ignores the safety period with a duration of Maximum Client Lead Time; a 'true' value has the effect that the DHCP server starts offering IP addresses from the partner's scope immediately after this transition, without waiting for existing leases allocated by the partner and not known by this system to time out. A 'true' value increases the risk that duplicate addresses are offered; if the transition to PARTNER-DOWN state is likely to be caused by a failure of the partner system rather than a communications problem, this risk is reduced.
Failover statistics for pool	The failover statistics for each pool
Dropped Invalid Packets	The number of BNDUPD packets that were dropped because the packet was malformed
Failover Shutdown	The number of BNDUPD packets that were dropped because the failover state if the DHCP server instance is shut down
Lease Already Expired	The number of BNDUPD packets that were dropped because the corresponding lease has expired
Maximum Lease Count Reached	The number of BNDUPD packets that were dropped because the maximum number of leases were reached
Subnet Not Found	The number of BNDUPD packets that were dropped because a valid subnet could not be found for the lease
Range Not Found	The number of BNDUPD packets that were dropped because a valid include range could not be found for the lease.
Host Conflict	The number of BNDUPD packets that were dropped because this DHCP server instance has already leased this address to another host
Address Conflict	The number of BNDUPD packets that were dropped because this DHCP server instance has already leased another address to this host
Peer Conflict	The number of BNDUPD packets that were dropped because the failover peer has leased an address within a subnet range of which the failover control is set to **local** on this local DHCP server instance

Table 60 **Failover Pool Statistics Field Descriptions (Continued)**

Field	Description
Persist Congestion	The number of BNDUPD packets that were dropped because of persistence congestion on this DHCP server instance
No Lease Hold Time Configured	The number of BNDUPD packets that were dropped because the lesae hold time is zero on this DHCP server instance
Lease Not Found	The number of Binding Database Update (BNDUPD) remove packets were dropped because the corresponding lease could not be found.
Number of pools found	The total number of pools found

pool-threshold-stats

Syntax **pool-threshold-stats** [*pool-name*] **detail** [**format** {**exact** | **scientific**}]
pool-threshold-stats [*pool-name*]

Context show>router>dhcp6>server

Description This commands displays pool level threshold stats of local DHCPv6 server. A minimum-free threshold needs to be configured before system collects threshold stats for the prefix.

The stats for each threshold are calculated based on the configured minimum-free prefix length.

For example, a /59 prefix is provision in the local DHCPv6 server, and the server allocated two PD leases, one /62 and one /63. And there is a /63 minimum threshold configured. So the threshold stats are calculated based on /63 as the base unit(block). So the value of “current used block” would be 3 because there is one /62 lease and one /63 lease, so it equals to total three /63.

Parameters *pool-name* — Specifies the name of the pool in local DHCPv6 server

detail — Displays detailed output

format — Specifies the format in the display to be either **exact** or **scientific**

Output The following output is an example of pool threshold statistics information.

Sample Output

```
show router 500 dhcp6 local-dhcp-server "d6" pool-threshold-stats "1"
===============================================================================
Server "d6"
===============================================================================
  Operational state         : inService
```

```
-------------------------------------------------------------------------------
Pool                        : 1
-------------------------------------------------------------------------------
  Stable leases             : 2
  Advertised leases         : 0
  ---------------------------------------------------------------------------
  Threshold   Used   Peak   Too low   Depleted   Peak timestamp
  ---------------------------------------------------------------------------
  /62         25%    25%    N         N          01/21/2015 21:52:12
  /63         19%    19%    N         N          01/21/2015 21:52:12
```

The command shown above displays an overview of pool level thresholds in the specified pool:

- The **Peak** field indicates the peak value of used
- The **Too low** field indicate if the configured minimum-free threshold is exceed
- The **Depleted** field indicate if there is no available prefix with the length in the provisioned prefix
- The **Peak timestamp** field indicates the time of peak used value

```
show router 500 dhcp6 local-dhcp-server "d6" pool-threshold-stats "1" detail
===============================================================================
Server "d6"
===============================================================================
  Operational state         : inService
-------------------------------------------------------------------------------
Pool                        : 1
-------------------------------------------------------------------------------
  Stable leases             : 2
  Advertised leases         : 0
  ---------------------------------------------------------------------------
  Threshold                 : /62
  ---------------------------------------------------------------------------
  Current Provisioned Blks  : 8.000000x10^0
  Current Used Blks         : 2.000000x10^0
  Current Free Blks         : 6.000000x10^0
  Current Used Percent      : 25%
  Current Used Peak Blks    : 2.000000x10^0
  Current Used Peak Percent : 25%
  Current Used Peak Time    : 01/21/2015 21:52:12
  Current Free Percent      : 75%
  Current Free Too Low      : N
  Current Free Depleted     : N
  Local Provisioned Blks    : 8.000000x10^0
  Local Used Blks           : 2.000000x10^0
  Local Free Blks           : 6.000000x10^0
  Local Used Peak Blks      : 2.000000x10^0
  Local Used Peak Percent   : 25%
  Local Used Peak Time      : 01/21/2015 21:52:12
  Remote Provisioned Blks   : 0.000000x10^0
  Remote Used Blks          : 0.000000x10^0
  Remote Free Blks          : 0.000000x10^0
  Remote Used Peak Blks     : 0.000000x10^0
  Remote Used Peak Percent  : 0%
  Remote Used Peak Time     : 01/21/2015 21:47:39
  Peak Reset Time           : 01/21/2015 21:47:39
```

```
Valid Data                : Y
-------------------------------------------------------------------------------
Threshold                 : /63
-------------------------------------------------------------------------------
Current Provisioned Blks  : 1.600000x10^1
Current Used Blks         : 3.000000x10^0
Current Free Blks         : 1.300000x10^1
Current Used Percent      : 19%
Current Used Peak Blks    : 3.000000x10^0
Current Used Peak Percent : 19%
Current Used Peak Time    : 01/21/2015 21:52:12
Current Free Percent      : 81%
Current Free Too Low      : N
Current Free Depleted     : N
Local Provisioned Blks    : 1.600000x10^1
Local Used Blks           : 3.000000x10^0
Local Free Blks           : 1.300000x10^1
Local Used Peak Blks      : 3.000000x10^0
Local Used Peak Percent   : 19%
Local Used Peak Time      : 01/21/2015 21:52:12
Remote Provisioned Blks   : 0.000000x10^0
Remote Used Blks          : 0.000000x10^0
Remote Free Blks          : 0.000000x10^0
Remote Used Peak Blks     : 0.000000x10^0
Remote Used Peak Percent  : 0%
Remote Used Peak Time     : 01/21/2015 21:47:39
Peak Reset Time           : 01/21/2015 21:47:39
Valid Data                : Y
```

The above command displays detailed statistics of all pool level thresholds in the specified pool:

- **Blks** in the output means the minimum free prefix length.
- **Valid Data** output indicates whether the data you see is valid or not. The data is invalid when a background stats update is scheduled or busy.

Table 61 describes pool threshold statistics output fields.

Table 61 **Pool Threshold Statistics Field Descriptions**

Field	Description
Operational state	The operational state of the local DHCP server instance unknown — The operational state is unknown inService — The operational state is in service outOfService — The operational state is out of service transition — The operational state is in transition waitPersistence — The DHCP server instance is waiting for a persistence action to complete.
Pool	The pool ID

Table 61 **Pool Threshold Statistics Field Descriptions (Continued)**

Field	Description
Stable Leases	The number of stable leases
Advertised leases	The number of advertised leases
Threshold	The prefix level threshold
Current Provisioned Blks	The number of provisioned blocks in this pool
Current Used Blks	The number of used blocks in this pool
Current Free Blks	The number of free blocks in this pool
Current Used Percent	The percentage of used prefixes with the minimum free threshold length in the pool compared to the number of provisioned prefixes
Current Used Peak Blks	A 64-bit word of the peak value of the number of used blocks in the pool with a prefix length
Current Used Peak Percent	The peak value of the number of used prefixes with the minimum free threshold length in the pool as a percentage of the provisioned prefixes. This depends on the current failover state of the DHCP server or pool
Current Used Peak Time	The time at which the peak value of the number of used prefixes in the pool was reached
Current Free Percent	The percentage of free prefixes with the minimum free threshold length in the pool compared to the number of provisioned prefixes. This depends on the current failover state of the DHCP server or pool
Current Free Too Low	The number of free prefixes with the minimum free threshold length available in the pool that is below the configured number of prefixes with the minimum free threshold length
Current Free Depleted	The number of prefixes with the minimum free threshold length available in the pool
Local Provisioned Blks	A 64-bit word of the number of provisioned prefixes with the minimum free threshold length local in the pool
Local Used Blks	A 64-bit word of the number of used prefixes with the minimum free threshold length local in the pool
Local Free Blks	A 64-bit word of the number of free prefixes with the minimum free threshold length local in the pool

Table 61 **Pool Threshold Statistics Field Descriptions (Continued)**

Field	Description
Local Used Peak Blks	A 64-bit word of the peak value of the number of used prefixes with the minimum free threshold length local in the pool
Local Used Peak Percent	The peak value of the number of used prefixes with the minimum free threshold length local in the pool as a percentage of the provisioned prefixes
Local Used Peak Time	The time at which the peak value of the number of used prefixes local in the pool was reached
Remote Provisioned Blks	A 64-bit word of the number of provisioned prefixes with the minimum free threshold length remote in the pool
Remote Used Blks	A 64-bit word of the number of used prefixes with the minimum free threshold length remote in the pool
Remote Free Blks	A 64-bit word of the number of free prefixes with the minimum free threshold length remote in the pool

prefix-ext-stats

Syntax **prefix-ext-stats** *ipv6-address/prefix-length*
prefix-ext-stats pool *pool-name*

Context show>router>dhcp6>server

Description This command displays extended statistics per DHCPv6 prefix in local DHCPv6 server.

The following statistics are included in output:

- The number of stable leases in the prefix
- The number of provisioned /64 address block in the prefix
- The number of used /64 address block in the prefix
- The number of free /64 address block in the prefix
- The percentage of used address (with /64 address block)
- The percentage of free address (with /64 address block)

For each statistic (except for Provisioned Addresses), there is current value and peak value, peak value is the highest value since prefix creation or last reset via command **clear>router** *rt-id*>**dhcp6**>**local-dhcp-server** *svr-name*>**prefix-ext-stats**.

When the **pool** parameter is used, the statistics of each prefix in the pool will be displayed.

Parameters *ipv6-address/prefix-length* — Specifies the IPv6 prefix

pool-name — The name of DHCPv6 local server pool

Output The following output is an example of prefix extended stats information.

Sample Output

```
show router 500 dhcp6 local-dhcp-server "d6" prefix-ext-stats 2001:ABCD::/62
===============================================================================
Extended statistics for prefix 2001:ABCD::/62
===============================================================================
                          Current         Peak            TimeStamp
-------------------------------------------------------------------------------
Local:
  Failover Oper State     Active
  Stable Leases           0               0               01/07/2013 19:54:52
  Provisioned Blks        4
  Used Blks               0               0               01/07/2013 19:54:52
  Free Blks               4               4               01/07/2013 19:54:52
  Used Pct                0               0               01/07/2013 19:54:52
  Free Pct                100             100             01/07/2013 19:54:52
Last Reset Time                                           01/07/2013 19:54:52
-------------------------------------------------------------------------------
Number of entries         1
===============================================================================
```

Table 62 describes extended prefix statistics output fields.

Table 62 **Prefix Extended Statistics Field Descriptions**

Field	Description
Current	The current extended prefix statistics
Peak	The peak value since the last reset
TimeStamp	The date and time of the last reset
Failover Oper State	The operational state of the DHCP specified prefix
Stable Leases	The number of stable leases
Provisioned Blks	The number of provisioned blocks
Used Blks	The number of used blocks
Free Blks	The number of free blocks
Used Pct	The percentage of extended prefixes in use
Free Pct	The percentage of free prefixes in use
Last Reset Time	The date and time of the last reset
Number of entries	The total number of entries

prefix-stats

Syntax **prefix-stats** *ipv6-address/prefix-length*
prefix-stats pool *pool-name*

Context show>router>dhcp6>local-dhcp-server

Description This command displays prefix statistics.

Parameters *ipv6-address* — Specifies the base IP address of the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

pool-name — Specifies the name of DHCPv6 local server pool

prefix-threshold-stats

Syntax **prefix-threshold-stats pool** *pool-name* **detail** [**format** {**exact** | **scientific**}]
prefix-threshold-stats pool *pool-name*
prefix-threshold-stats *ipv6-address*/*prefix-length* **detail** [**format** {**exact** | **scientific**}]
prefix-threshold-stats *ipv6-address*/*prefix-length*

Context show>router>dhcp6>server

Description This commands displays prefix level threshold stats of local DHCPv6 server prefix. A minimum-free threshold needs to be configured before system collects threshold stats for the prefix.

The stats for each threshold are calculated based on the configured minimum-free prefix length.

For example, a /59 prefix is provision in the local DHCPv6 server, and the server allocated two PD leases, one /62 and one /63. And there is a /63 minimum threshold configured. So the threshold stats are calculated based on /63 as the base unit(block). So the value of “current used block” would be 3 because there is one /62 lease and one /63 lease, so it equals to total three /63.

Parameters **pool** *pool-name* — Specifies the name of the pool in local DHCPv6 server, up to 32 characters

detail — Displays detailed output statistics

format — Specifies that the number format in the display will be either **exact** or **scientific**

ipv6-address/prefix-length — Specifies the IPv6 prefix with prefix length

Values

ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d

x [0 to FFFF]H
d [0 to 255]D
prefix-length [1 to 128]

Output The following output is an example of prefix threshold information.

Sample Output

```
show router 500 dhcp6 local-dhcp-server "d6" leases
===============================================================================
Leases for DHCPv6 server d6
===============================================================================
IP Address/Prefix                          Lease State       Remaining   Fail
  Link-local Address                                         LifeTime    Ctrl
-------------------------------------------------------------------------------
8888:0:0:ffe0:db8:/62
  fe80::3:ffff:fe00:db8:111                    stable            18h19m2s    local
8888:0:0:ffe4::/63
  fe80::3:ffff:fe00:db8:211                    stable            19h49m37s   local
-------------------------------------------------------------------------------
2 leases found
===============================================================================
show router 500 dhcp6 local-dhcp-server "d6" prefix-threshold-stats pool "1"
===============================================================================
Server "d6"
===============================================================================
  Operational state         : inService
-------------------------------------------------------------------------------
Pool                        : 1
-------------------------------------------------------------------------------
  Stable leases             : 2
  Advertised leases         : 0
-------------------------------------------------------------------------------
Prefix                      : 8888:0:0:ffe0:db8:/59
-------------------------------------------------------------------------------
  Stable leases             : 2
  Advertised leases         : 0
  Draining                  : N
  ---------------------------------------------------------------------------
  Threshold   Used   Peak   Too low   Depleted   Peak timestamp
  ---------------------------------------------------------------------------
  /62         25%    25%    Y         N          01/20/2015 23:51:36
  /63         19%    19%    N         N          01/21/2015 05:00:53
===============================================================================
```

The command shown above displays an overview of prefix level thresholds in the specified pool:

- The **Peak** field indicates the peak value of used.
- The **Too low** field indicate if the configured minimum-free threshold is exceed.
- The **Depleted** field indicate if there is no available prefix with the length in the provisioned prefix.
- The **Peak** timestamp field indicates the time of peak used value.

```
show router 500 dhcp6 local-dhcp-server "d6" prefix-threshold-stats pool "1" detail
===============================================================================
Server "d6"
===============================================================================
  Operational state         : inService
-------------------------------------------------------------------------------
Pool                        : 1
-------------------------------------------------------------------------------
  Stable leases             : 2
  Advertised leases         : 0
-------------------------------------------------------------------------------
Prefix                      : 8888:0:0:ffe0::/59
-------------------------------------------------------------------------------
  Stable leases             : 2
  Advertised leases         : 0
  Draining                  : N
  -----------------------------------------------------------------------------
  Threshold                 : /62
  -----------------------------------------------------------------------------
  Current Provisioned Blks  : 8.000000x10^0
  Current Used Blks         : 2.000000x10^0
  Current Free Blks         : 6.000000x10^0
  Current Used Percent      : 25%
  Current Used Peak Blks    : 2.000000x10^0
  Current Used Peak Percent : 25%
  Current Used Peak Time    : 01/21/2015 21:59:02
  Current Free Percent      : 75%
  Current Free Too Low      : N
  Current Free Depleted     : N
  Local Provisioned Blks    : 8.000000x10^0
  Local Used Blks           : 2.000000x10^0
  Local Free Blks           : 6.000000x10^0
  Local Used Peak Blks      : 2.000000x10^0
  Local Used Peak Percent   : 25%
  Local Used Peak Time      : 01/21/2015 21:59:02
  Remote Provisioned Blks   : 0.000000x10^0
  Remote Used Blks          : 0.000000x10^0
  Remote Free Blks          : 0.000000x10^0
  Remote Used Peak Blks     : 0.000000x10^0
  Remote Used Peak Percent  : 0%
  Remote Used Peak Time     : 01/21/2015 21:59:02
  Peak Reset Time           : 01/21/2015 21:59:02
  Valid Data                : Y
  -----------------------------------------------------------------------------
  Threshold                 : /63
  -----------------------------------------------------------------------------
  Current Provisioned Blks  : 1.600000x10^1
  Current Used Blks         : 3.000000x10^0
  Current Free Blks         : 1.300000x10^1
  Current Used Percent      : 19%
  Current Used Peak Blks    : 3.000000x10^0
  Current Used Peak Percent : 19%
  Current Used Peak Time    : 01/21/2015 21:59:13
  Current Free Percent      : 81%
  Current Free Too Low      : N
  Current Free Depleted     : N
  Local Provisioned Blks    : 1.600000x10^1
  Local Used Blks           : 3.000000x10^0
```

```
  Local Free Blks           : 1.300000x10^1
  Local Used Peak Blks      : 3.000000x10^0
  Local Used Peak Percent   : 19%
  Local Used Peak Time      : 01/21/2015 21:59:13
  Remote Provisioned Blks   : 0.000000x10^0
  Remote Used Blks          : 0.000000x10^0
  Remote Free Blks          : 0.000000x10^0
  Remote Used Peak Blks     : 0.000000x10^0
  Remote Used Peak Percent  : 0%
  Remote Used Peak Time     : 01/21/2015 21:59:13
  Peak Reset Time           : 01/21/2015 21:59:13
  Valid Data                : Y
```

The command shown above displays detailed statistics of all prefix level thresholds in the specified pool:

- **Blks** means the minimum free prefix length.
- **Valid Data** output indicates whether the data is or is not valid.The data is invalid when a background stats update is scheduled or busy.

```
show router 500 dhcp6 local-dhcp-server "d6" prefix-threshold-
stats 8888:0:0: ffe0:db8:/59
===============================================================================
Server "d6"
===============================================================================
  Operational state         : inService
-------------------------------------------------------------------------------
Pool                        : 1
-------------------------------------------------------------------------------
  Stable leases             : 2
  Advertised leases         : 0
-------------------------------------------------------------------------------
Prefix                      : 8888:0:0:ffe0:db8:/59
-------------------------------------------------------------------------------
  Stable leases             : 2
  Advertised leases         : 0
  Draining                  : N
  -----------------------------------------------------------------------------
  Threshold   Used   Peak   Too low   Depleted   Peak timestamp
  -----------------------------------------------------------------------------
  /62         25%    25%    N         N          01/21/2015 21:59:02
  /63         19%    19%    N         N          01/21/2015 21:59:13
```

The command shown above displays an overview of prefix level thresholds in the specified provision prefix.

```
show router 500 dhcp6 local-dhcp-server "d6" prefix-threshold-stats 8888:0:0:ffe0::/
59 detail
===============================================================================
Server "d6"
===============================================================================
  Operational state         : inService
-------------------------------------------------------------------------------
Pool                        : 1
-------------------------------------------------------------------------------
  Stable leases             : 2
```

```
  Advertised leases         : 0
-------------------------------------------------------------------------------
Prefix                      : 8888:0:0:ffe0:db8:/59
-------------------------------------------------------------------------------
  Stable leases             : 2
  Advertised leases         : 0
  Draining                  : N
  -----------------------------------------------------------------------------
  Threshold                 : /62
  -----------------------------------------------------------------------------
  Current Provisioned Blks  : 8.000000x10^0
  Current Used Blks         : 2.000000x10^0
  Current Free Blks         : 6.000000x10^0
  Current Used Percent      : 25%
  Current Used Peak Blks    : 2.000000x10^0
  Current Used Peak Percent : 25%
  Current Used Peak Time    : 01/21/2015 21:59:02
  Current Free Percent      : 75%
  Current Free Too Low      : N
  Current Free Depleted     : N
  Local Provisioned Blks    : 8.000000x10^0
  Local Used Blks           : 2.000000x10^0
  Local Free Blks           : 6.000000x10^0
  Local Used Peak Blks      : 2.000000x10^0
  Local Used Peak Percent   : 25%
  Local Used Peak Time      : 01/21/2015 21:59:02
  Remote Provisioned Blks   : 0.000000x10^0
  Remote Used Blks          : 0.000000x10^0
  Remote Free Blks          : 0.000000x10^0
  Remote Used Peak Blks     : 0.000000x10^0
  Remote Used Peak Percent  : 0%
  Remote Used Peak Time     : 01/21/2015 21:59:02
  Peak Reset Time           : 01/21/2015 21:59:02
  Valid Data                : Y
  -----------------------------------------------------------------------------
  Threshold                 : /63
  -----------------------------------------------------------------------------
  Current Provisioned Blks  : 1.600000x10^1
  Current Used Blks         : 3.000000x10^0
  Current Free Blks         : 1.300000x10^1
  Current Used Percent      : 19%
  Current Used Peak Blks    : 3.000000x10^0
  Current Used Peak Percent : 19%
  Current Used Peak Time    : 01/21/2015 21:59:13
  Current Free Percent      : 81%
  Current Free Too Low      : N
  Current Free Depleted     : N
  Local Provisioned Blks    : 1.600000x10^1
  Local Used Blks           : 3.000000x10^0
  Local Free Blks           : 1.300000x10^1
  Local Used Peak Blks      : 3.000000x10^0
  Local Used Peak Percent   : 19%
  Local Used Peak Time      : 01/21/2015 21:59:13
  Remote Provisioned Blks   : 0.000000x10^0
  Remote Used Blks          : 0.000000x10^0
  Remote Free Blks          : 0.000000x10^0
  Remote Used Peak Blks     : 0.000000x10^0
  Remote Used Peak Percent  : 0%
  Remote Used Peak Time     : 01/21/2015 21:59:13
```

```
Peak Reset Time           : 01/21/2015 21:59:13
Valid Data                : Y
```

The command displayed above displays detailed statistics of prefix level thresholds in the specified provision prefix.

Table 63 describes prefix threshold statistics output fields.

Table 63 **Prefix Threshold Statistics Field Descriptions**

Field	Description
Operational state	The operational state of the local DHCP server instance unknown — The operational state is unknown inService — The operational state is in service outOfService — The operational state is out of service transition — The operational state is in transition waitPersistence — The DHCP server instance is waiting for a persistence action to complete.
Pool	The number of pools in the specified prefix
Stable leases	The number of stable leases associated with the prefix
Advertised leases	The advertised leases associated with the prefix
Prefix	The specified prefix
Draining	The draining state Y — enabled N — disabled
Threshold	The prefix level threshold
Current Provisioned Blks	The number of provisioned blocks for this prefix
Current Used Blks	The number of used blocks for this prefix
Current Free Blks	The number of free blocks for this prefix
Current Used Percent	The percentage of used prefixes with the minimum free threshold length for this prefix
Current Used Peak Blks	A 64-bit word of the peak value of the number of used blocks for this prefix
Current Used Peak Percentage	The peak value of the number of used prefixes with the minimum free threshold length in the pool as a percentage of the provisioned prefixes.

Table 63 **Prefix Threshold Statistics Field Descriptions (Continued)**

Field	Description
Current Used Peak Time	The time at which the peak value of the number of used prefixes in the pool was reached
Current Free Percent	The percentage of free prefixes with the minimum free threshold length in the pool compared to the number of provisioned prefixes.
Current Free Too Low	The number of free prefixes with the minimum free threshold length available in the pool that is below the configured number of prefixes with the minimum free threshold length
Current Free Depleted	The number of prefixes with the minimum free threshold length available
Local Provisioned Blks	A 64-bit word of the number of provisioned prefixes with the minimum free threshold length local
Local Used Blks	The higher 64-bits word of the number of used prefixes with the minimum free threshold length
Local Free Blks	A 64-bit word of the number of free prefixes with the minimum free threshold length
Local Used Peak Blks	A 64-bit word of the peak value of the number of used prefixes with the minimum free threshold length
Local Used Peak Percent	The peak value of the number of used prefixes with the minimum free threshold length local in the pool as a percentage of the provisioned prefixes
Local Used Peak Time	The time at which the peak value of the number of used prefixes local in the pool was reached
Remote Provisioned Blks	A 64-bit word of the number of provisioned prefixes with the minimum free threshold length
Remote Used Blks	A 64-bit word of the number of used prefixes with the minimum free threshold length remote in the pool
Remote Free Blks	A 64-bit word of the number of free prefixes with the minimum free threshold length remote in the pool
Remote Used Peak Blks	The higher 64-bit word of the number of used prefixes with the minimum free threshold length
Remote Used Peak Percent	The peak value of the number of used prefixes with the minimum free threshold length remote in the pool as a percentage of the provisioned prefixes

Table 63 **Prefix Threshold Statistics Field Descriptions (Continued)**

Field	Description
Remote Used Peak Time	The time at which the peak value of the number of used prefixes remote in the pool was reached.
Peak Reset Time	The time at which the peak values have been reset
Valid Dada	The actual status of the threshold statistics data of the pool. Y — the data is up to date and may be used. N — the data is being recalculated in the background and is not stable for further use

lease-state

Syntax **lease-state** [**detail**] [**wholesaler** *service-id*] [**session** {**none**|**ipoe**|**ppp**}] [**router-advertisement-policy** *policy-name*]
lease-state [**detail**] **interface** *interface-name* [**wholesaler** *service-id*] [**session** {**none**|**ipoe**|**ppp**}] [**router-advertisement-policy** *policy-name*]
lease-state [**detail**] *ipv6-address ipv6-prefix*[/pre*fix-length*] [**wholesaler** *service-id*] [**session** {**none**|**ipoe**|**ppp**}] [**router-advertisement-policy** *policy-name*]
lease-state [**detail**] **mac** *ieee-address* [**wholesaler** *service-id*] [**session** {**none**|**ipoe**|**ppp**}] [**router-advertisement-policy** *policy-name*]

Context show>service>id>dhcp6

Description This command displays DHCPv6 lease state information.

Note: The **wholesaler** *service-id* parameter is applicable only in the VPRN context.

Parameters **detail** — Displays detailed lease state information

service-id — The service ID of the wholesaler. When specified in this context, SAP, SDP, interface, IP address and MAC parameters are ignored. This parameter only applies to the 7750 SR.

Values *service-id*: 1 to 2147483647

session — Shows DHCPv6 lease states for clients that are associated with an IPoE session or for clients that are associated with a PPP session or for clients that are not associated with an IPoE session

Values none, ipoe, ppp

policy-name — Specifies the name of the router advertisement policy

interface-name — Specifies the interface name, up to 32 characters

ipv6-address[*Imask*] — Shows information for the specified IPv6 address and mask

ieee-address — Shows information for the specified 48-bit MAC address, expressed in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff*, where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers

Output The following output is an example of DHCP6 lease state information.

Sample Output

```
*A:eng-BNG-2# show  service id 1000 dhcp6 lease-state detail
===============================================================================
DHCP lease states for service 1000
===============================================================================
Service ID           : 1000
IP Address           : 2001:1000:0:4::1/128
Client HW Address    : 00:00:10:10:12:12
Subscriber-interface : sub-int-01
Group-interface      : grp-int-01
SAP                  : 1/1/20:841
Termination Type     : local
Up Time              : 0d 00:00:13
Remaining Lease Time : 0d 23:59:47
Remaining SessionTime: N/A
Persistence Key      : N/A
Sub-Ident            : "host-1"
Sub-Profile-String   : "no-prof"
SLA-Profile-String   : "no-prof"
App-Profile-String   : ""
Lease ANCP-String    : ""
Lease Int Dest Id    : ""
Category-Map-Name    : ""
Dhcp6 ClientId (DUID): 00030001000010101212
Dhcp6 IAID           : 0
Dhcp6 IAID Type      : non-temporary
Dhcp6 Client Ip      : fe80:db8:200:10ff:fe10:1212
Primary-Dns          : N/A
Secondary-Dns        : N/A
Pool Name            : ""
Dhcp6 Server Addr    : ::10
Dhcp6 ServerId (DUID): 00030001d854ff000000
Dhcp6 InterfaceId    : test
Dhcp6 RemoteId       : N/A
Radius sub-if prefix : N/A
Router adv. policy   : ra-policy-01
Lease Info origin    : DHCP
ServerLeaseStart     : 02/13/2017 22:04:24
ServerLastRenew      : 02/13/2017 22:04:24
ServerLeaseEnd       : 02/14/2017 22:04:24
Session-Timeout      : N/A
IPoE|PPP session     : No
Radius User-Name     : "00:00:10:10:12:12"
-------------------------------------------------------------------------------
Number of lease states : 1
===============================================================================
```

```
*A:Dut-A#show service id 13 dhcp6 lease-state router-advertisement-policy ra-policy-
01
===============================================================================
DHCP lease state table, service 13
===============================================================================
IP Address      Mac Address       Sap/Sdp Id          Remaining  Lease    MC
                                                     LeaseTime  Origin   Stdby
-------------------------------------------------------------------------------
2013:bad:beaf::1/128
                00:00:00:00:00:13 1/1/1:13            23h47m34s  DHCP
2013:bad:beaf::2/128
                00:00:00:00:00:14 1/1/1:13            23h52m36s  DHCP
2013:bad:beaf::3/128
                00:00:00:00:00:15 1/1/1:13            23h54m17s  DHCP
-------------------------------------------------------------------------------
Number of lease states : 3
===============================================================================
*A:Dut-A#
```

Table 64 describes DHCP6 lease state router advertisement policy output fields.

Table 64 **DHCP6 Lease State Router Advertisement Policy Field Descriptions**

Field	Description
Service ID	The service ID provided by the access node to which this subnet is bound
IP Address	The IP address of the service
Client HW Address	The MAC address of the client
Subscriber-Interface	The subscriber interface for the service
Group-interface	The group interface for the service
Sap/Sdp ID	The service SAP and SDP IDs
Termination Type	The PPP session type
Up Time	The value of the up time at the time of the last modification of an entry
Remaining Lease Time	The time left for this lease
Remaining SessionTime	The time left for this session
Persistence Key	The persistence key
Sub-Ident	The subscriber identification policy name
Sub-Profile-String	The subscriber profile policy name

Table 64 **DHCP6 Lease State Router Advertisement Policy Field Descriptions (Continued)**

Field	Description
App-Profile-String	The application profile policy name
Lease ANCP-String	The ANCP string
Category-Map-Name	The category map name
Dhcp6 ClientID (DUID)	The DHCP client ID
Dhcp6 IAID	The DHCP6 Identity Association of the client
Dhcp6 IAID Type	The DHCP6 Identity Association of the client
Dhcp6 Client Ip	The DHCP6 IP address of the client
Primary-Dns	The primary DNS, if configured
Secondary-Dns	The secondary DNS, if configured
Pool Name	The pool name
Dhcp6 Server Addr	The DHCP6 server address
Dhcp6 ServerID (DUID)	The DHCP6 server ID
Dhcp6 InterfaceId	The DHCP6 interface ID
Dhcp6 RemoteId	The DHCP6 remote ID
Radius sub-if prefix	The RADIUS subscriber interface prefix
Router adv. policy	The router advertisement policy
Lease Info origin	The lease info origin
ServerLeaseStart	The date and time the lease was created
ServerLastRenew	The last renew was done for this lease
ServerLeaseEnd	The time left for this lease
Session-Timeout	The session timeout interval
IPoE|PPP	IPoE or PPP
Radius User-Name	The RADIUS user name
Number of lease states	The total number of lease states

statistics

Syntax **statistics** [**interface** *ip-int-name*]

Context show>router>dhcp6
show>service>id>dhcp6

Description This command displays statistics for DHCP6 relay and snooping.

Parameters *ip-int-name* — Specifies to display interface name information

Output The following output is an example of DHCP6 statistics information.

Sample Output

```
A:ALA-A# show router 1000 dhcp6 statistics
===============================================================================
DHCP6 Global Statistics (Service: 1000)
===============================================================================
Rx Packets                          : 16000
Tx Packets                          : 15041
Rx Malformed Packets                : 0
Rx Untrusted Packets                : 0
Client Packets Discarded            : 423
Client Packets Relayed              : 0
Client Packets Snooped              : 0
Client Packets Proxied (RADIUS)     : 0
Client Packets Proxied (Lease-Split) : 0
Server Packets Discarded            : 0
Server Packets Relayed              : 0
Server Packets Snooped              : 0
DHCP RELEASEs Spoofed               : 0
DHCP FORCERENEWs Spoofed            : 0
===============================================================================
A:ALA-A#
```

Table 65 describes DHCP6 output statistics.

Table 65 **DHCP6 Statistics Field Descriptions**

Field	Description
Rx Packets	The number of packets received
Tx Packets	The number of packets sent
Rx Malformed Packets	The number of malformed packets received from the DHCP6 clients
Rx Untrusted Packets	The number of untrusted packets received from the DHCP clients
Client Packets Discarded	The number of packets received from the DHCP clients that were discarded

Table 65 **DHCP6 Statistics Field Descriptions (Continued)**

Field	Description
Client Packets Relayed	The number of packets received from the DHCP clients that were forwarded
Client Packets Snooped	The number of packets received from the DHCP clients that were snooped
Client Packets Proxied (RADIUS)	The number of client packets proxied through RADIUS
Client Packets Proxied (Lease-Split)	The number of client packets proxied through a lease split
Server Packets Discarded	The number of discarded server packets
Server Packets Relayed	The number of relayed server packets
Server Packets Snooped	The number of snooped server packets
DHCP RELEASEs Spoofed	The number of spoofed DHCP releases
DHCP FORCERENEWs Spoofed	The number of spoofed FORCERENEW messages

summary

Syntax **summary** [**interface** *interface-name* | **saps**]

Context show>service>id>dhcp
show>service>id>dhcp6

Description This command displays DHCP or DHCP6 configuration summary information.

Parameters *interface-name* — Displays summary information for the specified existing interface

saps — Displays summary information for SAPs per interface

Output The following output is an example of a summary of DHCP information.

Sample Output

```
IES:
*A:ALA-48>show>service>id>dhcp# summary
===============================================================
DHCP Summary, service 700
===============================================================
Sap/Sdp                 Snoop  Used/       Info     Admin
                               Provided    Option   State
---------------------------------------------------------------
sap:1/1/9:0             No     0/0         Keep     Down
```

```
sap:1/1/25:0           No     0/0         Keep     Down
sdp:8:700              No     N/A         N/A      N/A
-------------------------------------------------------
Number of Entries : 3
-------------------------------------------------------
*A:ALA-48>show>service>id>dhcp#


VPLS:
*A:ALA-49>show>service# id 700 dhcp summary
===============================================================================
DHCP Summary, service 700
===============================================================================
Sap/Sdp                Snoop  Used/       Arp Reply   Info     Admin
                              Provided    Agent       Option   State
-------------------------------------------------------------------------------
sap:1/1/9:0            No     0/0         No          Keep     Down
sdp:2:222              No     N/A         N/A         N/A      N/A
sdp:2:700              No     N/A         N/A         N/A      N/A
-------------------------------------------------------------------------------
Number of Entries : 3
-------------------------------------------------------------------------------
*A:ALA-49>show>service#
```

The following example is for the 7750 SR:

```
VPRN:
*A:ALA-49>show>service# id 1 dhcp summary
===============================================================================================
DHCP Summary, service 1
===============================================================================================
Interface Name                   Arp      Used/                Info    Admin
  SapId/Sdp                      Populate Provided             Option  State
-----------------------------------------------------------------------------------------------
SpokeSDP                         No       0/0                  Keep    Down
  sdp:spoke-3:4                           0/0
test                             No       0/0                  Keep    Down
  sap:9/1/4:50/5                          0/0
to-ce1                           No       0/0                  Keep    Up
  sap:1/1/10:1                            0/0
-----------------------------------------------------------------------------------------------
Interfaces: 3
===============================================================================================
*A:ALA-49>show>service#

*A:ALA-48# show service id 88 dhcp summary saps
===============================================================================================
DHCP Summary, service 88
===============================================================================================
Interface Name                   Arp      Used/                Info    Admin
  SapId/Sdp                      Populate Provided             Option  State
-----------------------------------------------------------------------------------------------
SpokeTerm                        No       0/0                  Keep    Up
  sdp:spoke-3:3                           0/0
new-if                           No       0/1                  Keep    Up
  sap:1/2/19:0                            0/1
test123                          No       0/0                  Keep    Up
  sap:3/2/4:50/5                          0/0
```

```
testabc                          No       0/0                  Keep    Up
  sap:1/2/20:0                            0/0
-------------------------------------------------------------------------------
Interfaces: 4
===============================================================================
*A:ALA-48#


*A:ALA-48# show service id 88 dhcp summary interface SpokeTerm
===============================================================================
DHCP Summary, service 88
===============================================================================
Interface Name                   Arp      Used/                Info    Admin
  SapId/Sdp                      Populate Provided             Option  State
-------------------------------------------------------------------------------
SpokeTerm                        No       0/0                  Keep    Up
  sdp:spoke-3:3                           0/0
-------------------------------------------------------------------------------
Interfaces: 1
===============================================================================
*A:ALA-48#
```

Table 66 describes the DHCP summary output fields.

Table 66 **DHCP Summary Field Descriptions**

Field	Description
Interface Name	The DHCP interface name
Arp Populate	The state of ARP populate
Used/Provided	The number used or provided
Info Option	The info option
Admin State	The administrative state of the interface
Interfaces	The total number of interfaces

virtual-subnet

Syntax **virtual-subnet subscriber** *sub-ident*
virtual-subnet [**sap** *sap-id*]

Context show>service>id

Description This command displays currently recorded default gateway and subnets for all virtual subnets enabled for DHCPv4 hosts in the specified service.

Parameters *sub-ident* — Displays information relating to the specified subscriber ID

sap-id — Displays information relating to the specified SAP ID

Output The following output is an example of virtual subnet information.

Sample Output

```
show service id 500 virtual-subnet
===============================================================================
Virtual subnets in service 500
===============================================================================
Subscriber                  : 00:20:fc:1e:cd:52|1/1/9:200
-------------------------------------------------------------------------------
Default router              : 192.168.100.254
Subnet                      : 192.168.100.0/24
SAP                         : 1/1/9:200
-------------------------------------------------------------------------------
No. of subnets: 1
===============================================================================
```

Table 67 describes virtual subnet output fields.

Table 67 **Virtual Subnet Field Descriptions**

Field	Descriptions
Subscriber	The service identification that identifies the service in the service domain
Default router	The IP address of the default router received from either the DHCP or the RADIUS server
Subnet	The IP address of the subnet
SAP	The SAP associated with the service
No. of subnets	The total number of subnets for the service

3.6.2.1.8 Tools Commands

tools

Syntax **tools**

Context <root>

Description This command enables the context to enable useful tools for debugging purposes.

Parameters **dump** — Enables dump tools for the various protocols

perform — Enables tools to perform specific tasks

subscriber-mgmt

Syntax **subscriber-mgmt**

Context tools
tools>perform

Description This command enables tools to control subscriber management.

abort-startup-wait

Syntax **abort-startup-wait**

Context tools>perform>router>dhcp>server>failover
tools>perform>router>dhcp>server>pool
tools>perform>router>dhcp6>server
tools>perform>router>dhcp6>server>pool

Description This command aborts the running start-up-wait time, forcing the DCHP server to transition immediately to the COMMUNICATIONS-INTERRUPTED state.

The start-up-wait time is started after a boot to enable two redundant routers to synchronize the DHCP server lease database before transitioning to the (PRE-)NORMAL state when the last state record from the peer is synchronized. If the last state record is not received before the start-up-wait time expires, it is assumed that the remote DHCP server is not operational and the DHCP server goes in COMMUNICATIONS-INTERRUPTED state.

failover

Syntax **failover**

Context tools>perform>router>dhcp>server

tools>perform>router>dhcp>server>pool
tools>perform>router>dhcp6>server
tools>perform>router>dhcp6>server>pool

Description This command enables the context to perform local DHCP server failover tasks.

remap-lease-state

Syntax **remap-lease-state old-mac** *ieee-address* **mac** *ieee-address*
remap-lease-state sap *sap-id* [**mac** *ieee-address*]

Context tools>perform>subscr-mgmt

Description This command allows the remapping of all existing hosts if network card on CMTS/WAC side is changed is required.

When this command is executed, the following restrictions apply

- When **sap** is taken, all leases associated with the SAP are re-written.
  - For a SAP with a configured MAC in lease-populate command, this MAC will be taken.
  - For a SAP without a configured MAC the MAC from tools command will be taken.
  - For a SAP without a configured MAC and no MAC in tools command no action will be perform.
- When using the **old-mac** option, providing a new MAC *ieee-address* is mandatory.

This command is applicable only when dealing with DHCP lease states which were instantiated using l2header mode of DHCP operation.

Parameters **old-mac** *ieee-address* — Specifies the old MAC address to remap

mac *ieee-address* — Specifies that the provisioned MAC address will be used in the anti-spoofing entries for this SAP when the l2-header is enabled. This parameter may be changed mid-session. Existing sessions will not be re-programmed unless a **tools perform** command is issued for the lease.

sap-id — Specifies the physical port identifier portion of the SAP definition

When configured, the SAP parameter will remap all MAC addresses of DHCP lease states on the specified SAP. When no optional MAC parameter is specified, the **sap** *sap-id* command remaps all MAC addresses of lease states towards the MAC address specified in the l2-header configuration.

perform

Syntax perform

Context tools

Description This command enables the context to enable tools to perform specific tasks.

edit-ppp-session

Syntax **edit-ppp-session sap** *sap-id* [**user-name** *user-name*] [**mac** *mac-address*] [**session-id** *session-id*] [**ip** *ip-address*] [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*] [**inter-dest-id** *intermediate-destination-id*] [**ancp-string** *ancp-string*] [**app-profile-string** *app-profile-string*]

edit-ppp-session svc-id *service-id* [**user-name** *user-name*] [**mac** *mac-address*] [**session-id** *session-id*] [**ip** *ip-address*][**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*] [**app-profile-string** *app-profile-string*] [**inter-dest-id** *intermediate-destination-id*] [**ancp-string** *ancp-string*]

Context tools>perform>subscr-mgmt

Description This command modifies PPP session information.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition

ip-address — Specifies the IP address

sub-ident-string — Specifies a subscriber identification profile

sub-profile-string — Specifies the subscriber profile string, up to 16 characters

service-id — Specifies service identification number that identifies the service in the domain

intermediate-destination-id — Specifies the intermediate destination identifier which is encoded in the identification strings

ancp-string — Specifies the ASCII string of the DSLAM circuit ID name

app-profile-string — Specifies an application profile string

eval-lease-state

Syntax **eval-lease-state** [**svc-id** *service-id*] [**sap** *sap-id*] [**subscriber** *sub-ident-string*] [**ip** *ip-address*] [**mac** *mac-address*]

Context tools>perform>subscr-mgmt

Description This command evaluates lease state.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition

ip-address — Specifies the a server's IP address. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

sub-ident-string — Specifies the subscriber ID string, up to 32 characters

service-id — Specifies an existing service ID

Values 1 to 2147483647

local-user-db

Syntax **local-user-db** *local-user-db-name*

Context tools>perform>subscr-mgmt

Description This command enables tools for controlling the local user database.

Parameters *local-user-db-name* — Specifies the name of a local user database

ipoe

Syntax **ipoe**

Context tools>perform>subscr-mgmt>loc-user-db

Description This command contains the tools used to control IPoE entries in the local user database.

ipoe

Syntax **ipoe**

Context tools>perform>subscr-mgmt>loc-user-db

Description This command contains the tools used for controlling IPoE entries in the local user database.

authentication

Syntax **authentication user-name** *user-name* **password** *password*] [**circuit-id** *circuit-id*] [**circuit-id-hex** *circuit-id-hex*] [**mac** *ieee-address*] [**remote-id** *remote-id*] [**derived-id** *derived-id*] [**mac** *ieee-address*] [**remote-id-hex** *remote-id-hex*] [**sap-id** *sap-id*] [**service-name** *service-name*]

Context tools>perform>subscr-mgmt>loc-user-db>ppp

Description This command authenticates PPP user name. As local user database PAP/CHAP authentication can only be used when the local user database is connected to the PPP node under the group interface, the user lookup will be performed with match-list username.

Parameters *password* — Specifies the password of this host, up to 32 characters.

ieee-address — Specifies information about the MAC address of the PPP session.

remote-id — Specifies that information that goes into the remote-id sub-option in the DHCP relay packet, up to 255 characters

circuit-id — Specifies the circuit ID from the Option 82

circuit-id-hex — Specifies the circuit ID in hexadecimal format from the Option 82

Values 0x0 to 0xFFFFFFFF (maximum 254 hex nibbles)

user-name — Specifies the PPP user name

service-name — Specifies the service name

host-lookup

Syntax **host-lookup** [**mac** *ieee-address*] [**remote-id** *remote-id-ascii*] [**sap-id** *sap-id*] [**service-id** *service-id*] [**string** *vso-string*] [**system-id** *system-id*] [**option60** *option-60-ascii*] [**circuit-id** *circuit-id-ascii*] [**circuit-id-hex** *circuit-id-hex*] [**option60-hex** *option60-hex*] [**remote-id-hex** *remote-id-hex*] [**derived-id** *derived-id*] [**ip-prefix** *ip-prefix/ip-prefix-length*]

Context tools>perform>subscr-mgmt>loc-user-db>ipoe

Description This command performs a lookup in the local user database. This command looks up the host with the match-list configured in the local user database.

Parameters *ieee-address* — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

remote-id — Specifies the information which goes into the remote-id sub-option in the DHCP relay packet, up to 255 characters

sap-id — Specifies a SAP identifier to be used

service-id — Specifies an existing subscriber service ID

Values 1 to 2147483647

vso-string — Specifies a Vendor Specific Option (VSO) string

system-id — Specifies the system ID

Values up to 255 characters maximum

hex-string — Specifies the content of option 60 for this lookup

Values 0x0 to 0xFFFFFFFF (maximum 64 hex nibbles)

circuit-id — Specifies the circuit ID from the Option 82

circuit-id-hex — Specifies the circuit ID in hexadecimal format from the Option 82

Values 0x0 to 0xFFFFFFFF (maximum 254 hex nibbles)

sap-id — Specifies a SAP ID, up to 255 characters maximum

option-60-ascii — Specifies the Vendor-Identifying Vendor Option to match against. Option 60 is encoded as Type - Length - Value and must be formatted in ASCII

option-60-hex — Specifies the Vendor-Identifying Vendor Option to match against. Option 60 is encoded as Type - Length - Value and must be formatted in HEX

ppp

Syntax **ppp**

Context tools>perform>subscr-mgmt>loc-user-db

Description This command contains the tools used to control PPP entries in the local user database.

host-lookup

Syntax **host-lookup** [**circuit-id** *circuit-id*] [**circuit-id-hex** *circuit-id-hex*] [**derived-id** *derived-id*] [**mac** *ieee-address*] [**remote-id** *remote-id*] [**remote-id-hex** *remote-id-hex*] [**sap-id** *sap-id*] [**service-name** *service-name*] [**user-name** *user-name*]

Context tools>perform>subscr-mgmt>loc-user-db>ppp

Description This command performs a lookup in the local user database.

Parameters *ieee-address* — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

remote-id — Specifies the information that goes into the remote-id sub-option in the DHCP relay packet, up to 255 characters

user-name — Specifies a user name, up to 128 characters

service-name — Specifies a PPP service name, up to 255 characters

circuit-id circuit-id — Specifies the circuit ID from the Option 82

circuit-id-hex — Specifies the circuit ID in hexadecimal format from the Option 82

Values 0x0 to 0xFFFFFFFF (maximum 254 hex nibbles)

derived-id *derived-id* — Specifies an ASCII string that uniquely identifies a host, and is derived by a Python script from packet content available during a DHCP transaction

pool

Syntax **pool** *pool-name*

Context tools>perform>router>dhcp>server

Description This command performs local DHCP server IP address pool tasks.

Parameters *pool-name* — Specifies the pool name

create-sticky-lease

Syntax **create-sticky-lease** *host-name* [**mac** *ieee-address*] [**circuit-id** *circuit-id*] [**client-id** *client-id*] [**requested-ip-address** *ip-address*] [**circuit-id-hex** *circuit-id-hex-string*] [**client-id-hex** *client-id-hex-string*]

Context tools>perform>router>dhcp>server>pool

Description This command tries to create a lease-state in the specified local-dhcp-server pool with the supplied parameters. The hostname will uniquely identify this lease for subsequent operations. mac/circuit-id and client-id should also have a unique combination of values according to the configured user-ident variables.

If the requested-ip-address is provided the server will try to create a lease with this address. If this address is not available in the pool, the command will fail.

Parameters *host-name* — Specifies a string to identify the lease, up to 32 characters. If not specified, it can be internally generated based on MAC@ or other key.

ieee-address — Specifies the MAC address in xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx format

circuit-id — Specifies the circuit-id string, up to 255 characters

client-id — Specifies the client-id string, up to 255 characters

ip-address — Specifies the IPv4 address in dotted notation a.b.c.d

dhcp-rx-stats

Syntax **dhcp-rx-stats** [**clear**]

Context tools>dump

Description This command displays the DHCP received packet statistics.

Parameters **clear** — Clears DHCP RX statistics

Output The following output is an example of dhcp rx stats information.

Sample Output

```
*A:# tools dump dhcp-rx-stats
===============================================================================
DHCP Received Packet Statistics
===============================================================================
     Type                  Received     Forwarded       Dropped  Dropped(ESM)
```

```
-------------------------------------------------------------------------------
IPv4 DISCOVER                    72            72             0             0
     OFFER                        2             2             0             0
     REQUEST                      4             4             0             0
     DECLINE                      0             0             0             0
     ACK                          1             1             0             0
     NAK                          0             0             0             0
     RELEASE                      0             0             0             0
     INFORM                       0             0             0             0
     FORCERENEW                   0             0             0             0
     LEASEQUERY                   0             0             0             0
     LEASEUNASSIGNED              0             0             0             0
     LEASEUNKNOWN                 0             0             0             0
     LEASEACTIVE                  0             0             0             0
     RENEW                        2             2             0             0
-------------------------------------------------------------------------------
IPv6 SOLICIT                     0             0             0             0
     ADVERTISE                   0             0             0             0
     REQUEST                     0             0             0             0
     CONFIRM                     0             0             0             0
     RENEW                       0             0             0             0
     REBIND                      0             0             0             0
     REPLY                       0             0             0             0
     RELEASE                     0             0             0             0
     DECLINE                     0             0             0             0
     RECONFIGURE                 0             0             0             0
     INFO_REQUEST                0             0             0             0
     RELAY_FORW                  0             0             0             0
     RELAY_REPLY                 0             0             0             0
     LEASEQUERY                  0             0             0             0
     LEASEQUERY_REPLY            0             0             0             0
-------------------------------------------------------------------------------
     Total                      81            81             0             0
-------------------------------------------------------------------------------
Maximum queue length          : 0
Maximum outst pbufs total     : 1
Maximum outst pbufs to client : 1
===============================================================================
```

force-partner-down

Syntax **force-partner-down**

Context tools>perform>router>dhcp>server>failover
tools>perform>router>dhcp>server>pool
tools>perform>router>dhcp6>server
tools>perform>router>dhcp6>server>pool

Description This command forces a DHCP server to transition from the COMMINICATIONS-INTERRUPTED state to the PARTNER-DOWN state.

It is important to note that this command should only be used if the remote DHCP server is out of service, else duplicate address allocations could occur.

In case of a communication failure between two redundant DHCP servers, the DHCP server transitions to the COMMUNICATIONS-INTERRUPTED state. In this state, the DHCP server allocates IP addresses for new leases from the local and address-driven address-ranges only. This prevents duplicate address allocations.

If during the partner-down-delay time, the communication with the remote DHCP server is not restored, then the DHCP server transitions to the PARTNER-DOWN state. In this state, the DHCP server also allocates IP addresses for new leases from the remote ranges. A DHCP server should only reach this state when the remote DHCP server is out of service, else IP address conflicts can occur.

pool

Syntax **pool** *pool-name*

Context tools>perform>router>dhcp>server

Description This command performs local DHCP server IP address pool tasks.

Parameters *pool-name* — Specifies the pool name, up to 32 characters

send-force-renew

Syntax **send-force-renew** *ip-address*

Context tools>perform>router>dhcp>server

Description This command sends a DHCP FORCERENEW message on behalf of the DHCP server to the DHCP client with specified IP address. A DHCP lease with the specified IP address must exist in the DHCP server.

Parameters *ip-address* — Specifies the DHCP client's IP address in dotted notation a.b.c.d

3.6.2.1.9 Clear Commands

dhcp

Syntax **dhcp**

Context clear>router

Description This command enables the context to clear and reset DHCP entities.

local-dhcp-server

Syntax **local-dhcp-server** *server-name*

Context clear>router>dhcp
clear>router>dhcp6

Description This command clears DHCP server data.

Parameters *server-name* — Clears data for the specified local DHCP server

declined-addresses

Syntax **declined-addresses** *ip-address*[*/mask*]
declined-addresses pool *pool-name*

Context clear>router>dhcp>local-dhcp-server

Description This command clears declined DHCP addresses.

Parameters *pool-name* — Clears information about the declined pool name

ip-address[/mask] — Clears information about the declined IP address and mask

failover-pool-stats

Syntax **failover-pool-stats** [*pool-name*]

Context clear>router>dhcp>local-dhcp-server
clear>router>dhcp6>local-dhcp-server

Description This command clears failover pool statistics.

Parameters *pool-name* — Clears information about the pool name which is defined in the system

leases

Syntax **leases** *ip-address*[*/mask*] [*state*]
leases all [*state*]

Context clear>router>dhcp>local-dhcp-server
clear>router>dhcp6>local-dhcp-server

Description This command clears DHCP leases.

Parameters *ip-address*[*/mask*] — Clears the specified IP address and mask

state — Clears the state of the lease to be removed

Values offered, stable, force-renew-pending, remove-pending, held, internal, internal-orphan, internal-held, sticky

pool-ext-stats

Syntax **pool-ext-stats** [*pool-name*]

Context clear>router>dhcp>local-dhcp-server
clear>router>dhcp6>local-dhcp-server

Description This command clears extended pool statistics.

Parameters *pool-name* — Clears information about the pool name

server-stats

Syntax **server-stats**

Context clear>router>dhcp>local-dhcp-server

Description This command clears all server statistics.

sticky-leases

Syntax **sticky-leaseshostname** [*hostname*]
sticky-leases hostname-prefix [*hostname-prefix*]

Context clear>router>dhcp>local-dhcp-server

Description This command clears sticky leases.

Parameters *hostname* — Clears information about the name of the host for which an entry is cleared

hostname-prefix — Clears information about the sub-string of all hostnames starting with that sub-string

subnet-ext-stats

Syntax **subnet-ext-stats** *ip-address*[/*mask*]
subnet-ext-stats pool *pool-name*

Context clear>router>dhcp>local-dhcp-server

Description This command clears extended subnet statistics.

Parameters *ip-address*[/*mask*] — Clears the specified IP address and mask

pool-name — Clears statistics for the specified pool name

statistics

Syntax **statistics** [*ip-int-name* | *ip-address*]

Context clear>router>dhcp

Description This command clears DHCP statistics.

Parameters *ip-int-name* — Clears DHCP statistics for the specified interface name

ip-address — Clears DHCP statistics for the specified IP address

dhcp6

Syntax **dhcp6**

Context clear>router

Description This command enables the context to clear and reset DHCP6 entities.

pool-threshold-stats

Syntax **pool-threshold-stats** [*pool-name*]

Context clear>router>dhcp6>server

Description This commands resets the peak stats in the pool level threshold stats in the specified pool. If the pool name is not specified, then the peak stats in all pools in the server will be reset.

Parameters *pool-name* — Clears information about the specified local DHCPv6 server poo

prefix-ext-stats

Syntax **prefix-ext-stats** *ipv6-address/prefix-length*
prefix-ext-stats pool *pool-name*

Context clear>router>dhcp6>server

Description This command reset the begin time of peak values in output of the **show router** *rt-id* **dhcp6 local-dhcp-server** *svr-name* **prefix-ext-stats** command/

Parameters *ipv6-address/prefix-length* — Clears information about the specified IPv6 prefix

pool-name — Clears information about the specified local DHCPv6 server pool name

prefix-threshold-stats

Syntax **prefix-threshold-stats** *ipv6-address/prefix-length*
prefix-threshold-stats pool *pool-name*

Context clear>router>dhcp6>server

Description This commands resets the peak stats in the prefix level threshold stats in the specified provision prefix or pool.

Parameters *pool-name* — Clears information about the specified pool in local DHCPv6 server

ipv6-address/prefix-length — Clears information about the specified IPv6 prefix with prefix length

server-stats

Syntax **server-stats**

Context clear>router>dhcp6>server

Description This command reset all stats of the specified local DHCPv6 server.

statistics

Syntax **statistics**

Context clear>router>dhcp6

Description This command clears DHCP6 statistics.

statistics

Syntax **statistics** [**interface** *ip-int-name* | *ipv6-address*]

Context clear>service>id>dhcp6

Description This command clears DHCP statistics.

Parameters *ip-int-name* — Clears DHCP statistics for the specified interface name

ip-address — Clears DHCP statistics for the specified IP address

statistics

Syntax **statistics** [**interface** *ip-int-name* | *ip-address*]

Context clear>router>dhcp

Description This command clears DHCP statistics.

Parameters *ip-int-name* — Clears all ARP cache entries for the IP interface with the specified name

ip-address — Clears all ARP cache entries for the specified IP interface with the specified IP address

lease-state

Syntax **lease-state** [**no-dhcp-release**]
lease-state [**port** *port-id*] [**inter-dest-id** *intermediate-destination-id*] [**no-dhcp-release**]
lease-state [**port** *port-id*] **no-inter-dest-id** [**no-dhcp-release**]
lease-state ip-address *ip-address*[*/mask*] [**no-dhcp-release**]
lease-state mac *ieee-address* **no-dhcp-release**
lease-state sap *sap-id* [**no-dhcp-release**]
lease-state sdp *sdp-id:vc-id* [**no-dhcp-release**]

Context clear>service>id>dhcp

Description This command clears DHCP lease state information.

Parameters **no-dhcp-release** — Specifies that the node will clear the state without sending the DHCP release message

port-id — The DHCP lease state local specifies that the DHCP lease state is learned by either a SAP or SDP. When the value is SAP, the value indicates the SAP for which this entry contains information.

intermediate-destination-id — Specifies the intermediate destination identifier, up to 32 characters, which is encoded in the identification strings

ip-address[/mask] — Clears information about the specified IPv6 address and mask

ieee-address — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

sap-id — Clears the specified lease state SAP information

sdp-id — Clears information about the specified SDP ID

Values 1 to 17407

vc-id — Clears information about the specified virtual circuit ID on the SDP ID

Values 1 to 4294967295

lease-state

Syntax **lease-state all** [**no-dhcp-release**]
lease-state ipv6-address *ipv6-prefix*[*/prefix-length*] [**no-dhcp-release**]
lease-state mac *ieee-address* [**no-dhcp-release**]
lease-state sap *sap-id* [**no-dhcp-release**]

Context clear>service>id>dhcp6

Description This command clears DHCP6 lease state information.

Parameters **all** — Clears all lease-states.

no-dhcp-release — Specifies that the node will clear the state without sending the DHCP release message

ipv6-prefix[/prefix-length] — Clears routes only matching the specified ip-address and length and only applies to the 7750 SR

Values		
ipv6	ipv6-prefix[/pref*:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
		x:x:x:x:x:x:d.d.d.d
		x: [0 to FFFF]H
		d: [0 to 255]D
	prefix-length:	1 to 128

ieee-address — Clears DHCP6 MAC address lease state information. The 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

sap-id — Clears the specified lease state SAP information

lease-state

Syntax **lease-state** [**no-dhcp-release**]
lease-state [**port** *port-id*] [**inter-dest-id** *intermediate-destination-id*] [**no-dhcp-release**]
lease-state [**port** *port-id*] **no-inter-dest-id** [**no-dhcp-release**]
lease-state ip-address *ip-address*[/*mask*] [**no-dhcp-release**]
lease-state mac *ieee-address* [**no-dhcp-release**]
lease-state sap *sap-id* [**no-dhcp-release**]
lease-state sdp *sdp-id*[:*vc-id* [**no-dhcp-release**]

Context clear>service>id>dhcp

Description This command clears DHCP lease state information for this service.

Parameters **no-dhcp-release** — Specifies that the node clears the state without sending the DHCP release message

port-id — Clears DHCP lease state local information that the DHCP lease state is learned by either a SAP or SDP. When the value is SAP, the value indicates the SAP for which this entry contains information.

ip-address — Clears information about the specified IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that is used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

ieee-address — Clears information about the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

intermediate-destination-id — Clears information about the intermediate destination identifier, up to 32 characters, which is encoded in the identification strings.

sap-id — Clears the specified lease state SAP information

sdp-id — Clears information about the SDP ID

Values 1 to 17407

vc-id — The virtual circuit ID on the SDP ID to be cleared

Values 1 to 4294967295

leases

Syntax **leases** [*ipv6-address/prefix-length*] [*type*] [*state*]
leases all [*type*] [*state*]

Context clear>router>dhcp6>server

Description This command removes the specified leases in the specified local DHCPv6 server.

Parameters *ipv6-address/prefix-length* — Specifies the prefix of the leases to be removed

type — Specifies the type of the lease to be removed

Values pd, slaac, wan-host

state — Specifies the state of the lease to be removed

Values advertised, remove-pending, held, internal, internal-orphan, internal-offered

all — Removes all leases of specified type or state

3.6.2.1.10 Debug Commands

dhcp

Syntax [**no**] **dhcp** [**interface** *ip-int-name*]
[**no**] **dhcp mac** *ieee-address*
[**no**] **dhcp sap** *sap-id*

Context debug>router>ip

Description This command enables DHCP debugging.

The **no** form of this command disables debugging.

Parameters *ip-int-name* — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

ieee-address — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

sap-id — Specifies the physical port identifier portion of the SAP definition

detail-level

Syntax **detail-level** {**low** | **medium** | **high**}
no detail-level

Context debug>router>ip>dhcp
debug>router>ip>dhcp6
debug>router>local-dhcp-server

Description This command debugs the DHCP tracing detail level.

Parameters **low** — Displays a low detail level for DHCP debugging

medium — Displays a medium detail level for DHCP debugging

high — Displays a high detail level for DHCP debugging

local-dhcp-server

Syntax [**no**] **local-dhcp-server** *server-name* [**lease-address** *ip-prefix*[*/prefix-length*]]
[**no**] **local-dhcp-server** *server-name* [**mac** *ieee-address*]
[**no**] **local-dhcp-server** *server-name* [**link-local-address** *ipv6z-address*]

Context debug>router

Description This command enables, disables or configures debugging for a local DHCP server.

Parameters *server-name* — Specifies an existing local DHCP server name

ip-prefix[/prefix-length] — Specifies the IP prefix and prefix length of the subnet

Values ip-prefix a.b.c.d (host bits must be 0)
length 0 to 32

ieee-address — Specifies that the provisioned MAC address for the local DHCP server

ipv6z-address —

ipv6-address:	x:x:x:x:x:x:x:x [-interface]
	x:x:x:x:x:x:d.d.d.d [-interface]
	x: [0 to FFFF]H
	d: [0 to 255]D
interface	up to 32 characters, mandatory for link local addresses

mode

Syntax **mode** {**dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}
no mode

Context debug>router>ip>dhcp
debug>router>ip>dhcp6
debug>router>local-dhcp-server

Description This command debugs the DHCP tracing detail level.

Parameters **dropped-only** — Only displays dropped packets

ingr-and-dropped — Only displays ingress packet and dropped packets

egr-ingr-and-dropped — Displays ingress, egress and dropped packets

dhcp6

Syntax **dhcp6** [*ip-int-name*]
no dhcp6

Context debug>router>ip

Description This command enables DHCP debugging

The **no** form of this command disables debugging.

Parameters *ip-int-name* — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

4 Stateless Address Auto-configuration (SLAAC)

Note: The information in this section applies only to the 7750 SR.

4.1 SLAAC Management Principles

In a Triple Play network, client devices can use SLAAC to dynamically obtain their IP address and other network configuration information.

1. During bootup, the client sends a Router Solicit (RS) message to get an IP prefix.
2. The BNG address server can assign a prefix statically to the subscriber through RADIUS or LUDB. Or, dynamically through the use of the local address server.
3. The BNG address server replies to the client with a Router Advertisement which contains a /64 prefix.

4.2 Configuration Overview

The ICMP6 Router Solicit is the primary trigger for SLAAC host creation. It is also possible to use the DHCPv4 message to trigger a SLAAC host creation using the “IPoE-linking” feature. The SLAAC host can use RADIUS or LUDB authentication, as well as bypass authentication. Address assignment can be assigned statically or dynamically. Static prefix assignment is accomplished through RADIUS or LUDB. Dynamic prefix assignment requires the use of the local-address-server (reusing the local DHCPv6 server), and a pool name returned from RADIUS or LUDB. The DHCPv6 server for SLAAC is used for address management only, there are no lease state associated with SLAAC users. The DHCPv6 server can be shared with regular DHCPv6 users as well.

4.3 Router-Solicit Trigger

The following example shows a router-solicit triggered configuration.

```
*A:eng-BNG-2>config>service>vprn>sub-if>grp-if>ipv6# info
----------------------------------------------
                        router-solicit
                            no shutdown
                        exit
```

To add authentication to the above configuration, there are two options.

- For RADIUS authentication, similar to DHCP and PPP authentication, add a RADIUS policy under the group interface.
- For LUDB, add the following to the **router-solicit** configuration.

```
*A:eng-BNG-2>config>service>vprn>sub-if>grp-if>ipv6# info
----------------------------------------------
                        router-solicit
                            user-db "slaac-users"
                            no shutdown
                        exit
```

4.4 SLAAC Address Assignment

After an RS is received to trigger the creation of a SLAAC host, address assignment can be provided statically or dynamically.

4.5 Static SLAAC Prefix Assignment

If using RADIUS, the attribute “framed-ipv6-prefix” VSA is used. The attribute must use a /64 prefix.

```
*A:eng-BNG-2>config>subscr-mgmt>loc-user-db>ipoe>host# info
----------------------------------------------
                    ipv6-slaac-prefix 2001::/64
```

4.6 Dynamic SLAAC Prefix Assignment

SLAAC prefix can be dynamically assigned to a user at real time. Prefixes are assignment through the local DHCPv6 pool. Therefore a DHCPv6 pool must be defined first. The following displays an example configuration.

```
*A:eng-BNG-2>config>service>vprn>dhcp6# info
----------------------------------------------
       local-dhcp-server "dhcp6-server" create
              use-pool-from-client
              pool "pool-01" create
                     prefix 2001::/32 wan-host create
                     exit
              exit
        exit
```

To associate the DHCPv6 server for SLAAC address assignment, the following configuration is used. Notice the server name configured under **local-address-assignment dhcp6-server** matches the name configured under the DHCPv6 pool.

```
*A:eng-BNG-2>config>service>vprn>sub-if>grp-if# info
----------------------------------------------
                    local-address-assignment
                        ipv6
                            client-application ppp-slaac ipoe-slaac
                            server "dhcp6-server"
                        exit
                        no shutdown
                    exit
```

In order to specify the pool to be used for SLAAC prefix assignment, the pool name can either be returned from LUDB or RADIUS.

If using RADIUS, the attribute “Alc-slaac-ipv6-pool” is used.

If using LUDB, the following configuration is used.

```
*A:eng-BNG-2>config>subscr-mgmt>loc-user-db>ipoe>host# info
----------------------------------------------
                    ipv6-slaac-prefix-pool "pool-01"
```

In this example, the pool named “pool-01” provisioned in the LUDB or returned from RADIUS will match the pool name configured in the DHCP6 server. A prefix from the 2001::/32 pool will be assigned to the SLAAC subscribers.

4.7 SLAAC Prefix Replacement

An SLAAC host prefix can be replaced with the VSA Alc-Ipv6-Slaac-Replacement-Prefix. This VSA is only supported through CoA or through the **tools>subscr-mgmt>coa** command. When a CoA is triggered, the original SLAAC host session terminates from the BNG. Depending on the accounting mode, an accounting stop message may be sent. Immediately following the termination of the original SLAAC host session, an SLAAC host with the new replacement prefix is created on the system. The SLAAC host inherits all of the original host attributes, such as the subscriber profile and SLA profile. Nokia recommends against combining the VSA Alc-Ipv6-Slaac-Replacement-Prefix with other VSAs. An error in any of the VSAs can cause the SLAAC host recreation to fail.Service can only be restored after the subscriber performs an address request and authenticates.

The RA prefix replacement generates a single router advertisement containing both the old and new prefixes. The old prefix will have both the valid and preferred lifetime parameters set to 0, informing the subscriber to deprecate the prefix as soon as possible. The new prefix will have the valid and preferred lifetime parameters set as per the operator configuration. The subscriber can continue to use the old prefix for up to two hours after the RA. During the two hours, the 7750 SR drops subscriber traffic that does not match the anti-spoof criteria.

This feature is supported on an MCS setup. For persistence, the replacement SLAAC prefix is stored as the subscriber new prefix.

SLAAC prefixes that were assigned through local address assignment cannot be replaced. This feature ensures that when SLAAC replacement is performed, the address origin is not changed. This feature only replaces SLAAC host prefixes and cannot, for example, replace a DHCPv6 host with an SLAAC prefix. This feature is not supported for PPPoE sessions.

5 Point-to-Point Protocol over Ethernet Management

5.1 PPPoE

Note: The information in this section applies only to the 7750 SR.

A Broadband Remote Access Server (BRAS) is a device that terminates PPPoE sessions. The Point-to-Point Protocol (PPP) is used for communications between a client and a server. Point-to-Point Protocol over Ethernet (PPPoE) is a network protocol used to encapsulate PPP frames inside Ethernet frames.

Ethernet networks are packet-based, unaware of connections or circuits. Using PPPoE, Nokia users can dial from one router to another over an Ethernet network, then establish a point-to-point connection and transport data packets over the connection. In this application subscriber hosts can connect to the router using a PPPoE tunnel. There are two command available under PPPoE to limit the number of PPPoE hosts, one to set a limit that is applied on each SAP of the group-interface and one to set the limit per group-interface.

PPPoE is commonly used in subscriber DSL networks to provide point-to-point connectivity to subscriber clients running the PPP protocol encapsulated in Ethernet. IP packets are tunneled over PPP using Ethernet ports to provide the client's software or RG the ability to dial into the provider network. Most DSL networks were built with the use of PPPoE clients as a natural upgrade path from using PPP over dial-up connections. Because the PPP packets were used, many of the client software was reusable while enhancements were made such that the client could use an Ethernet port in a similar manner as it did a serial port. The protocol is defined by RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*.

PPPoE has two phases, the discovery phase and the session phase.

- Discovery: The client identifies the available servers. To complete the phase the client and server must communicate a session-id. During the discovery phase all packets are delivered to the PPPoE control plane (CPM or MDA). The IOM identifies these packets by their ethertype (0x8863).

  - PPPoE Active Discovery Initiation (PADI). This broadcast packet is used by the client to search for an active server (Access Concentrator) providing access to a service.
  - PPPoE Active Discovery Offer (PADO): If the access server can provide the service it should respond with a unicast PADO to signal the client it may request connectivity. Multiple servers may respond and the client may choose a server to connect to.
  - PPPoE Active Discovery Request (PADR): After the client receives a PADO it uses this unicast packet to connect to a server and request service.
  - PPPoE Active Discovery Session-confirmation (PADS) A server may respond to the client with this unicast packet to establish the session and provide the session-id. Once the PADS was provided the PPP phase begins.

- Session: Once the session ID is established connectivity is available for the duration of the session, using ethertype 0x8864. Either client or server can terminate a session.

  During the life of the session the packets may be uniquely identified by the client's MAC address and session-id. The session can terminate either by PADT sent by the client or server or by an LCP Terminate-Request packet.

During session creation, the following occurs:

- PADI (control packet upstream): This packet is delivered to the control plane. The control plane checks the service tag for service name. In the case multiple nodes are in the same broadcast domain the service tag can be used to decide whether to respond to the client. A relay tag can also be present.
- PADO (control packet downstream): The packet is generated by the control plane as response to the PADI message. The packet is forwarded to the client using the unicast packet directed at the client's MAC address. The node populates the AC-name tag and service tag. The packet sources the forwarding Ethernet MAC address of the node. In the case SRRP is used on the interface, it uses the gateway address as the source MAC. When in a backup state, the packet is not generated.
- PADR (control packet upstream): This packet is delivered to the control plane. The packet is destined to the node's MAC address. The control plane then generates the PADS to create the session for this request.
- PADS: (control packet downstream): The control plane prepares for the session creation and sends it to the client using the client's MAC address. The session-id (16-bit value) is unique per client. The session-id is populated in the response. Once a session-id is generated, the client uses it in all packets. In cases that the server does not agree with the client's populated service tags, the PADS can be used to send a service error tag with a zero session-id to indicate the failure.

- PADT (control packet upstream/downstream): The packet is used to terminate a session. It can be generated by either the control plane or the client. The session-id must be populated. The packet is a unicast packet.
- PPP session creation supports the LCP authentication phase.

During a session, the following forwarding actions occur:

- Upstream, in the PPPoE before PPP phase, there is no anti-spoofing. All packets are sent to the CPM. During anti-spoof lookup with IP and MAC addressing, regular filtering, QoS and routing in context continue. All unicast packets are destined to the node's MAC address. Only control packets (broadcast) are sent to the control plane. Keep-alive packets are handled by the CPM.
- Downstream, packets are matched in the subscriber lookup table. The subscriber information provides queue and filter resources. The subscriber information also provides PPPoE information, such as the dest-mac-address and session-id, to build the packet sent to the client.

PPPoE-capable interfaces can be created in a subscriber interface in both IES and VPRN services (VPRN is supported on the 7750 SR only). Each SAP can support one or more PPPoE sessions depending on the configuration. A SAP can simultaneously have static hosts, DHCP leases and PPPoE sessions. The number of PPPoE sessions is limited per SAP, SLA profile or subscriber profile.

RADIUS can be used for authentication. IP addresses can be provided by both RADIUS and the local IP pool, with the possibility of choosing the IP pool through RADIUS.

DHCP clients and PPPoE clients are allowed on a single SAP or group interface. If DHCP clients are not allowed, the operator should not enable lease-populate and similarly if PPPoE clients are not allowed, the operator should not enable the PPPoE node.

Note: The DHCP node can be enabled when only PPPoE clients are allowed since the DHCP relay function can be used for IP retrieval.

The DHCP lease-populate is for DHCP leases only. A similar command host-limit is made available under PPPoE for limits on the number of PPPoE hosts. The existing per sla-profile instance host limit is for combined DHCP and PPPoE hosts for that instance.

- For authentication, local and RADIUS are supported.
  - RADIUS is supported through an existing policy. A username attribute has been added.

  - For PAP/CHAP, a local user database is supported and must be referenced from the interface configuration.
- The host configuration can come directly from the local user database or from the RADIUS or DHCP server. A local host configuration is allowed through a local DHCP server with a local user database.
- IP information can be obtained from the local user database, RADIUS, a DHCP server, or a local DHCP server.

If IP information is returned from a DHCP server. PPPoE options such as the DNS name are retrieved from the DHCP ACK and provided to the PPPoE client. An open authentication option is maintained for compatibility with existing DHCP-based infrastructure.

The DHCP server can be configured to run on a loopback address with a relay defined in the subscriber or group interfaces. The DHCP proxy functionality that is provided by the DHCP relay (getting information from RADIUS, lease-split, option 82 rewriting) cannot be used for requests for PPPoE clients.

5.1.1 PPPoE Authentication and Authorization

5.1.1.1 General Flow

When a new PPPoE session is setup, the authentication policy assigned to the group interface is examined to determine how the session should be authenticated.

If no authentication policy is assigned to the group interface or the **pppoe-access-method** is set to **none**, the local user database assigned to the PPPoE node under the group interface is queried either during the PADI phase or during the LCP authentication phase, depending on whether the match-list of the local user database contains the requirement to match on username. If the match-list does not contain the username option, PADI authentication is performed and can specify an authentication policy in the local user database host for an extra RADIUS PAP-CHAP authentication point.

If an authentication policy is assigned and the pppoe-access-method is set to PADI, the RADIUS server is queried for authenticating the session based on the information available when the PADI packet is received (any PPP user name and password are not known here). When it is set to PAP-CHAP, the RADIUS server is queried during the LCP authentication phase and the PPP user name and password is used for authentication instead of the user name and password configured in the authentication policy.

If this authentication is successful, the data returned by RADIUS or the local user database is examined. If no IP address was returned, the DHCP server is now queried for an IP address and possibly other information, such as other DHCP options and ESM strings.

The final step consists of complementing the available information with configured default values (ESM data), after which the host is created if sufficient information is available to instantiate it in subscriber management (at least subscriber ID, subscriber profile, SLA profile, and IP address).

The information that needs to be gathered is divided in three groups, subscriber ID, ESM strings, and IP data. Once one of the data sources has offered data for one of these groups, the other sources are no longer allowed to overwrite this data (except for the default ESM data). For example, if RADIUS provides an SLA profile but no subscriber ID and IP address, the data coming from the DHCP server (either through Python or directly from the DHCP option) can no longer overwrite any ESM string, only the subscriber ID and IP data. However, after the DHCP data is processed, a configured default subscriber profile is added to the data before instantiating the host.

5.1.1.2 RADIUS

Refer the 7750 SR and VSR RADIUS Attributes Reference Guide for attributes that are applicable in RADIUS authentication of a PPPoE session.

5.1.1.3 Local User Database Directly Assigned to PPPoE Node

The following are relevant settings for a local user database directly assigned to PPPoE node:

- Host identification parameters (user name only)
- Username
- Password
- Address
- DNS servers (under DHCP options)
- NBNS servers (under DHCP options)
- Identification (ESM) strings

Incoming PPPoE connections are always authenticated through the PPPoE tree in the local user database.

The match list for a local user database that is assigned directly to the PPPoE node under the group interface is always **user-name**, independent of the match list setting in the database.

For user-name matching, the incoming user name (user[@domain]) is always first converted to a user and a domain entity by splitting it on the first @-sign. If the no-domain parameter to the user name is given, the user component should be equal to the given user name, if the domain-only portion of the user name is given, the domain entity should be equal to the given user name and if no extra parameters are given, the user and domain components are concatenated again and compared to the given user name.

The option number for the identification strings is not used if the local user database is assigned directly to the PPPoE node (it is only necessary if it is connected to a local DHCP server). Any valid value may be chosen in this case (if omitted, the default value chosen is 254).

If a pool name is given for the address, this pool name is sent to the DHCP server in a vendor-specific sub-option of Option 82 to indicate from which pool the server should take the address. If the gi-address option is given for the address, this is interpreted as if no address was given.

5.1.1.4 Subscriber per PPPoE Session Index

The system keeps track of the number of PPPoE sessions active on a given SAP and assign a per SAP session index to each such that always the lowest free index is assigned to the next active PPPoE session. When PAP/CHAP RADIUS authentication is used, the PPPoE SAP session index can be sent to, and received from, the RADIUS server using the following VSA:

```
ATTRIBUTE Alc-SAP-Session-Index           180      integer
```

This is supported for all PPPoE sessions, including those using LAC and LNS, but is not supported in a dual-homing topology. It should only be used in a subscriber per VLAN model as the session index is per SAP.

The SAP session index allows PPPoE sessions to have their own set of queues for QoS and accounting purposes when using the same SLA profile name as that received from a RADIUS server. An example of this with multiple levels of HQoS egress scheduling is shown in Figure 31. Alternatively, this can be achieved by configuring per-session SPI sharing in the SLA profile as described in SLA Profile Instance Sharing.

Figure 31 **Egress QoS per PPPoE Session**

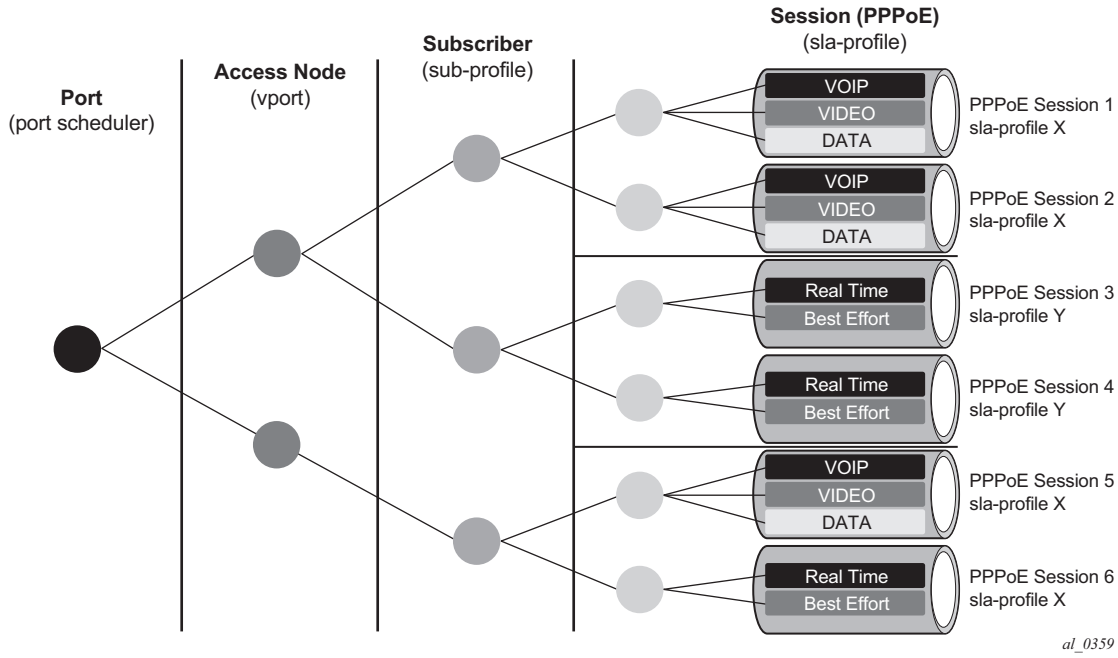

This requires a set of identical SLA profiles to be configured which only differ by an index being, for example, appended to their name. The SAP session index must be sent to RADIUS in the Access-Request message, which is achieved by configuring the RADIUS authentication policy to include it as follows:

```
configure subscriber-mgmt authentication-policy name
    include-radius-attribute
        [no] sap-session-index
```

The RADIUS server must then reflect the SAP session index back to the system in the RADIUS Access-Accept message together with the SLA profile name.

A Python script processes the RADIUS Access-Accept message to append the SAP session index to the SLA profile name to create the unique SLA profile name, in this example with the format:

sla-profile *sla-profile-name*.**suffix**

The exact format (for example, the separator used) is not fixed and just needs to match the pre-provisioned SLA profiles, while not exceeding 16 characters. This ensures that each PPPoE session is given its own SLA profile and consequently its own set of queues.

This processing is shown in Figure 32.

Figure 32 **Per PPPoE Session SLA Profile Selection**

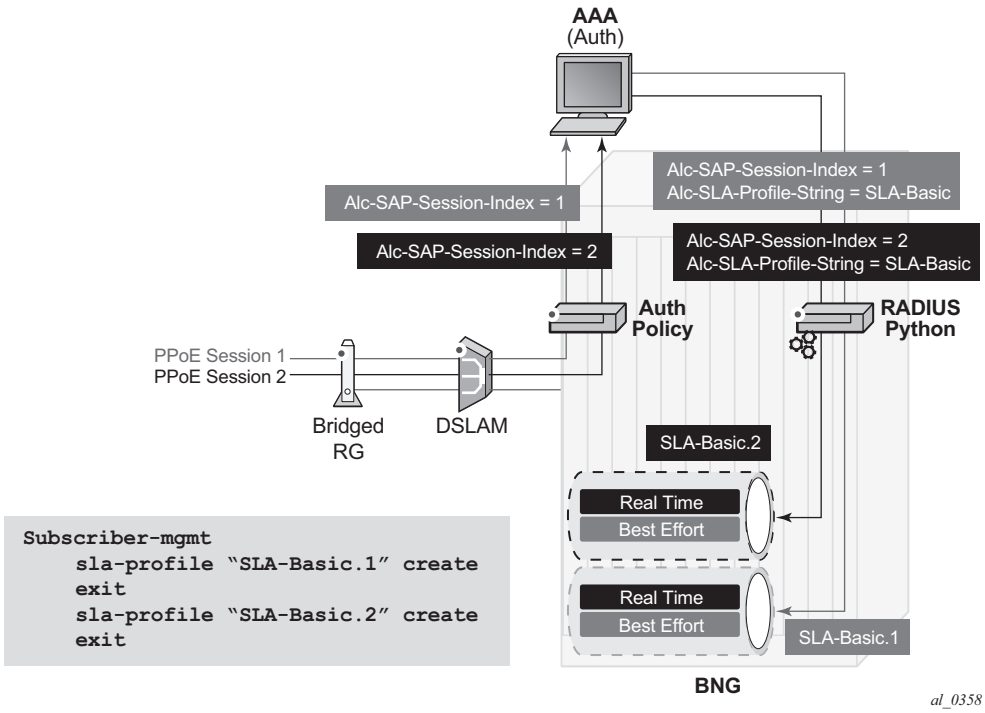

Below is an example Python script for this purpose:

```
import alc
import struct
from alc import radius
from alc import sub_svc

PROXY_STATE = 33
ALU = 6527
SLA_PROF_STR = 13
SAP_SESSION_INDEX = 180


##################################
## QoS for Multiple PPPoE Sessions
#   This script checks if a sap-session-
index (sid) is included in the authentication
#   accept. If present, the sla-profile-string (sla) is adapted to "sla.sid"

if alc.radius.attributes.isVSASet(ALU,SLA_PROF_STR):
    sla = alc.radius.attributes.getVSA(ALU,SLA_PROF_STR)
    if alc.radius.attributes.isVSASet(ALU,SAP_SESSION_INDEX):
        ssi = alc.radius.attributes.getVSA(ALU,SAP_SESSION_INDEX)
        suffix = "" .join(["%x" % ord(x) for x in ssi])
        alc.radius.attributes.setVSA(ALU,SLA_PROF_STR,sla + '.' + "%d" %
int(suffix,16))
```

In order to use a COA to change the SLA profile used, the new SLA profile name must be constructed with the same suffix (in this example) as that used for the current SLA profile. This is necessary in order to ensure unique use of a given provisioned SLA profile. This mandates that the SAP session index is included in the COA information. Two options are proposed to achieve this:

- The COA can specify a new SLA profile name and include the SAP session index. A Python script would then process the COA and construct the new SLA profile name to be used by appending the suffix in the same way as was done with the RADIUS Access-Accept.
- The COA could be using a RADIUS proxy which might make the first option unattractive. An alternative solution would be to use a Python script to append the suffix to the acct-session-id in all messages sent so that the suffix can be identified when a COA is received that uses the acct-session-id(+suffix) for session identification. This would need to be performed for all messages sent that include the acct-session-id. COAs would reference the session using the acct-session-id+suffix. A Python script would be required to remove the suffix and append it to the new SLA profile name. All messages received with the acct-session-id+suffix would be processed by the Python script to remove the suffix before sending the acct-session-id to the system.

  In order to ensure that the acct-session-id sent in RADIUS accounting messages is updated with the suffix, the user must configure include-radius-attribute sla-profile in the RADIUS accounting policy to be applied. The Python script needs to remove the suffix from the SLA profile and add it to the acct-session-id for all messages sent. Clearly the acct-session-id used by any external server would then be different to that seen on the system.

5.1.1.5 Local DHCP Server with Local User Database

If a DHCP server is queried for IP or ESM information, the following information is sent in the DHCP request:

- Option 82 sub-options 1 and 2 (Circuit-ID and Remote-ID): These options contain the Circuit-ID and Remote-ID that were inserted as tags in the PPPoE packets).
- Option 82 sub-option 9 vendor-id 6527 VSO 6 (client type): this value is set to 1 to indicate that this information is requested for a PPPoE client. The local DHCP server uses this information to match a host in the list of PPPoE users and not in the DHCP list.
- Option 82 sub-option 6 (Subscriber-ID): This option contains the user name that was used if PAP/CHAP authentication was performed on the PPPoE connection.

- Option 82 sub-option 13 (DHCP pool): This option indicates to the DHCP server that the address from the client should be taken from the specified pool. The local DHCP server only honors this request if the **use-pool-from-client** option is given in the server configuration.
- Option 82 sub-option 14 (PPPoE Service-Name): This option contains the service name that was used in the PADI packet during PPPoE setup.
- Option 60 (Vendor class ID): This option contains the string “ALU7XXXSBM” to identify the DHCP client vendor.
- The WT-101 access loop options are not sent for PPPoE clients

Local user database settings relevant to PPPoE hosts when their information is retrieved from the local DHCP server using this database:

- Host identification parameters (including user name)
- Address
- DNS servers (under DHCP options)
- NBNS servers (under DHCP options)
- Identification (ESM) strings

For user name matching, the incoming user name (user[@domain]) is always first converted to a user and a domain entity by splitting it on the first @-sign. If the no-domain parameter to the user name is given, the user component should be equal to the given user name, if the domain-only portion of the user name is given, the domain entity should be equal to the given user name and if no extra parameters are given, the user and domain components are concatenated again and compared to the given user name.

To prevent load problems, if DHCP lease times of less than 10 minutes are returned, these are not accepted by the PPPoE server.

5.1.2 Multiple Sessions Per MAC Address

To support MAC-concentrating network equipment, which translates the original MAC address of a number of subscribers to a single MAC address towards the PPPoE service, the 77x0 supports at most 1023 PPPoE sessions per MAC address. Each of these sessions are identified by a unique combination of MAC address and PPPoE session ID.

To set up multiple sessions per MAC, the following limits should be set sufficiently high:

- Maximum sessions per MAC in the PPPoE policy

- The PPPoE interface session limit in the PPPoE node of the group interface
- The PPPoE SAP session limit in the PPPoE node of the group interface
- The multiple-subscriber-sap limit in the subscriber management node of the SAP

If host information is retrieved from a local DHCP server, care must be taken that, although a host can be identified by MAC address, circuit ID, remote ID or user name, a lease in this server is only indexed by MAC address and circuit ID. For example, multiple sessions per MAC address are only supported in this scenario if every host with the same MAC address has a unique Circuit-ID value.

5.1.3 PPP Session Re-establishment

The **re-establish-session** command allows a host to re-establish a PPP session if the previous PPP session has yet to be terminated. The only way for a host to reconnect is to wait for the health check to fail or a manual termination by the operator. To allow a faster reconnect, the feature allows a PADI request to terminate the host previous session and allow the host to re-establish a new PPP session. As the old PPP session terminates, the accounting record also stops. The new PPP session starts a new accounting record; this is to ensure that the subscriber is not be charged for the unused time in the previous PPP session.

Note:

- Subscribers that use credit-control-policy along with PPP re-establish-session can experience a longer attempt at re-establishing a PPP connection.
- When PPP subscribers closes an L2TP connection, the terminate cause on the accounting record would show “user-request”. For cases, where PPP re-establishment is enabled, stale connections are closed without requests. The terminate cause on the accounting record would show “loss of service”.

5.1.4 Private Retail Subnets

IPoE and PPPoE are commonly used in residential networks and have expanded into business applications.

Both IPoE and PPPoE subscriber hosts terminate in a retail VPRN. It is possible for the subscriber to connect to one or more 7750 SRs for dual homing purposes. When the subscriber is dual-homed, routing between the two BNGs is required and can be performed with either a direct spoke SDP between the two nodes or by MPBGP route learning.

For PPPoE, both PADI and PAP/CHAP authentication are supported. The PPPoE session is negotiated with the parameters defined by the retail VPRN interface. Because the IP address space of the sub-mgmt host may overlap between VPRN services, the node must anti-spoof the packets at access ingress with the session ID.

When the **config**>**service**>**vprn**>**sub-if**>**private-retail-subnets** command is enabled on the subscriber interface, the retail IP subnets are not pushed into the wholesale context. This allows IP addresses to overlap between different retail VPRNs (those VPRNs can have IPoE and/or PPPoE sessions). If an operator requires both residential and business services, two VPRNs connected to the same wholesaler can be created and use the flag in only one of them.

To update a subscriber attribute (such as an SLA-profile or subscriber-profile), either perform a RADIUS CoA or enter the **tools**>**perform**>**subscriber-mgmt**>**coa** command. To identify a subscriber for a CoA, the subscriber IP address can be used as part of the subscriber's key. Because it is possible for different subscribers to use the same IP address in multiple private retail VPRNs, additional parameters are required in addition to the subscriber IP address. When performing a CoA on a retail subscriber for PPPoE hosts using a private retail subnet, the following conditions apply:

- If NAS-port-ID is used, the VSA Alc-Retail-Serv-Id must be included. The subscriber IP address and prefix must also be included.
- If Alc-serv-ID is used, the subscriber retail service ID must be referenced. The subscriber IP address and prefix must also be included.

When performing a CoA on a retail subscriber for IPoE hosts using a private retail subnet, the following conditions apply:

- If NAS-port-ID is used, the VSA Alc-Client-Hardware-Addr referencing the subscriber MAC must be included. The VSA Alc-Retail-Serv-Id must not be included. The subscriber IP address and prefix must also be included.
- If Alc-serv-ID is used, it must reference the subscriber retail wholesale service ID. The VSA Alc-Client-Hardware-Address and the subscriber IP address and prefix must also be included.

To perform a DHCP force renew on IPoE hosts, enter the **tools**>**perform**>**subscriber-mgmt**>**forcerenew** command. To identify between subscribers within a private retail VPRN where overlapping IP addresses are possible, the MAC address must be used instead of the IP address in the **tools** command.

Private retail subnet are supported on both numbered and unnumbered subscriber interfaces. RADIUS host creation is possible for IPoEv4 hosts.

When IPoE and PPPoE session terminates in the retail VPRN, the node must learn the retail VPRN service ID. This can be provided by the LUDB or RADIUS. If the local user database is used, the host configuration provides a reference to the VPRN service ID. If RADIUS is used, RADIUS can return the service ID VSA. The retail subscriber interface must also reference the subscriber interface on the wholesale subscriber interfaces.

The following are not supported on private retail subnets:

- ARP hosts
- static hosts
- MCS redundancy

ESM multicast is not supported for IPoE hosts that use overlapping IP address between private retail VPRNs. ESM multicast is only supported for IPoE hosts that have unique IP addresses in the system. It is possible for a IPoE subscriber to contain both type of hosts, only the hosts with overlapping private IP address will not support ESM multicast.

5.1.5 IPCP Subnet Negotiation

This feature enables negotiation between Broadband Network Gateway (BNG) and customer premises equipment (CPE) so that CPE is allocated both ip-address and associated subnet.

Some CPEs use the network up-link in PPPoE mode and perform dhcp-server function for all ports on the LAN side. Instead of wasting one subnet for P2P uplink, CPEs use allocated subnet for LAN portion as shown in Figure 33.

Figure 33 **CPEs Network Up-link Mode**

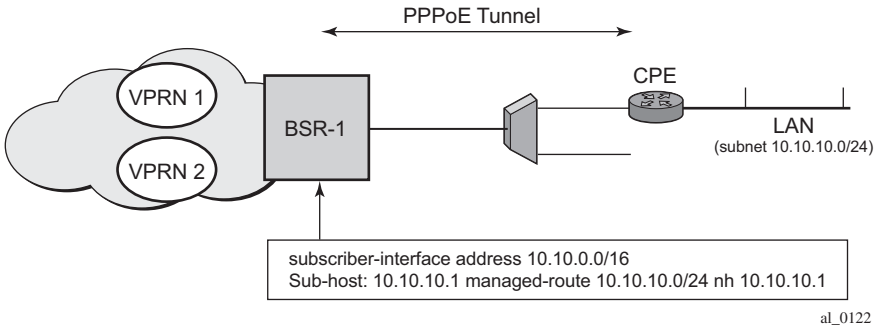

From a BNG perspective, the given PPPoE host is allocated a subnet (instead of / 32) by RADIUS, external dhcp-server, or local-user-db. And locally, the host is associated with managed-route. This managed-route is a subset of the subscriber-interface subnet, and also, subscriber-host ip-address is from managed-route range. The negotiation between BNG and CPE allows CPE to be allocated both ip-address and associated subnet.

5.1.5.1 Numbered WAN Support for Layer 3 RGs

Numbered WAN interfaces on RGs is useful to manage v6-capable RGs. Dual-stack RGs can be managed via IPv4. However, with v6-only RGs or with dual-stack RGs with private only v4 address, RGs require a globally routable v6 WAN prefix (or address) for management. This feature provides support to assign WAN prefix to PPP based Layer 3 RG using SLAAC. The feature also adds a new RADIUS VSA (Alc-PPP-Force-IPv6CP, Boolean) to control triggering of IPv6CP on completion of PPP LCP. RA messages are sent as soon as IPv6CP goes into open state, and the restriction to hold off on sending RAs until DHCP6-PD is complete in case of dual-stack PPP is no longer applicable.

5.1.6 IES as Retail Service for PPPoE Host

In this application, the PPPoE subscriber host terminates in a retail IES, the IES service ID can be obtained by either of following:

- Alc-Retail-Serv-Id attribute in radius access-accept packet.
- or retail-service-id config in local-user-db pppoe host if local user DB is used.

If MSAP is used then the SAP is created in the wholesale VPRN.

The PPPoE session is negotiated with the parameters defined by the wholesale VPRN group interface. The connectivity to the retailer is performed using the linkage between the two interfaces.

Due to the nature of IES service, there is no IP address overlap between different IES retails services, thus, the private-retail-subnets flag is not needed in this case.

5.1.7 Unnumbered PPPoX

Note: Unnumbered subscriber-interfaces are supported only for PPPoE, PPPoA and PPPoEoA (v4 and v6) hosts.

Unlike regular IP routes which are mainly concerned with next-hop information, subscriber-hosts are associated with an extensive set of parameters related to filtering, qos, stateful state (PPPoE/DHCP), antispoofing, and so on. Forwarding Database (FDB) is not suitable to maintain all this information. Instead, each subscriber host record is maintained in separate set of subscriber-host tables.

By pre-provisioning the IP prefix (IPv4 and IPv6) under the subscriber-interface and subscriber-interface>ipv6 CLI hierarchy, only a single prefix aggregating the subscriber host entries is installed in the FDB. This FDB entry points to the corresponding subscriber-host tables that contain subscriber-host records.

In case that IPv4/IPv6 prefix is not pre-provisioned, or the subscriber-hosts falls out of pre-provisioned prefix, each subscriber-host is installed in the FDB. The result of the subscriber-host FDB lookup points to the corresponding subscriber-host record in the subscriber-host table. This scenario is referred to as **unnumbered subscriber-interfaces**.

Unnumbered does not mean that the subscriber hosts do not have an IP address or prefix assigned. It only means that the IP address range out of which the address or prefix is assigned to the host does not have to be known in advance via configuration under the **subscriber-interface** or **subscriber-interface>ipv6** node.

An IPv6 example would be:

```
configure
    router/service
        subscriber-interface <name>
            ipv6
                [no] allow-unmatching-prefixes
                delegated-prefix-length <bits>
                subscriber-prefixes
```

This CLI indicates the following:

- There is no need for any indication of anticipated IPv6 prefixes in CLI.
- However, the **delegated-prefix-length** (DPL) command is required. The DPL (or the length of the prefix assigned to each Residential Gateway) must be known in advance. All Residential Gateways (or subscribers) under the same **subscriber-interface** share this pre-configured DPL.
- The DPL range is 48 — 64.
- If the prefix length in the received PD (via DHCP server, RADIUS or LUDB) and the DPL do not match, the host creation fails.
- If the assigned IP prefix/address (DHCPServer, RADIUS, LUDB) for the host falls outside of the CLI defined prefixes and the **allow-unmatching-prefixes** command is configured, then the new address/prefix automatically installs in the FDB.

5.1.8 Selective Backhaul of PPPoE Traffic using an Epipe Service

The router supports the redirection of PPPoE packets on ingress to a Layer 3 static subscriber SAP and PW-SAPs (for example, bound to an IES or VPRN service) in the upstream direction towards an Epipe service. The Epipe, which is bound to a specific SAP on a group interface of a Layer 3 service, is used to backhaul the PPPoE traffic towards a remote destination, such as, a wholesale service provider. Other non-PPPoE packets are still forwarded to the group interface on the IES or VPRN.

There is a one-to-one mapping from backhaul Epipe to the subscriber. A single backhaul Epipe service per subscriber SAP is supported. Only static configuration of subscriber SAPs is supported.

In the downstream direction, all traffic arriving on the backhaul Epipe is forwarded to the subscriber SAP and merged with IPoE traffic.

This architecture is shown in Figure 34.

Figure 34 **Redirection of Traffic to Epipe for Backhaul**

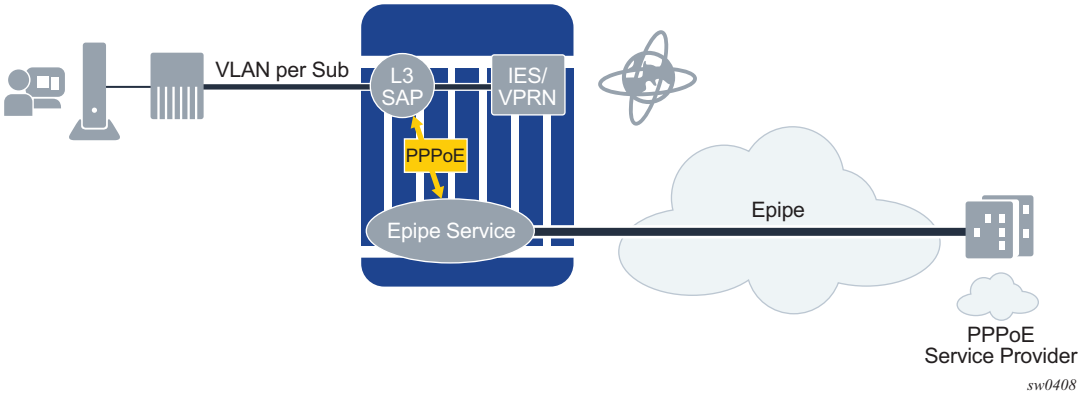

If the SAP on the Layer 3 service is configured to indicate PPPoE redirection, in addition to ESM anti-spoofing, then the following processing occurs in the upstream direction:

- All PPPoE traffic, such as traffic with Ethertype 0x8863 and 0x8864, skips the anti-spoof lookup and is then be forward to a configured Epipe service.
- All non-PPPoE is subject to anti-spoofing lookup. If the anti-spoofing fails, then the packet is dropped. Otherwise, processing continues using the host configuration.

The **fwd-wholesale** context in the Layer 3 SAP is used to configure the forwarding of PPPoE packets towards a specified Epipe service:

```
config
   service
      ies | vrpn
         subscriber-interface
            group-interface
               sap 1/1/1:10.20
                  anti-spoof ...
                  fwd-wholesale
                     pppoe 10
                     exit
```

The PPPoE option specifies that only packets matching Ethertype 0x8863 and 0x8864 are redirected to the Epipe of service ID '10'. This includes all PPPoE control plane packets. Note that the service IDs specified under **fwd-wholesale** cannot refer to a vc-switching Epipe service.

When **fwd-wholesale** is configured to an Epipe with a specified service ID, the system ensures that the Epipe exists and meets all the requirements to participate in the PPPoE redirect. There is no need for additional configuration under the Epipe. For the previous example, only the following configuration is required for the Epipe:

```
epipe 10 name "10" customer 1 create
   service-mtu 1400
   spoke-sdp 10:9 create
      no shutdown
      exit
   no shutdown
   exit
```

The system generates a CLI error if a user tries to configure additional SAPs on an Epipe that is already referenced from a **fwd-wholesale** context.

In the upstream direction, subscriber queues are used for IPoE packets destined for a local host and PPPoE backhaul traffic. However, CPM traffic is not consuming subscriber queue resources; QoS profiles are instantiated while **single-sub-parameters profiled-traffic-only** is enabled.

In the downstream direction, PPPoE traffic arriving on the Epipe is merged back into the subscriber SAP referencing the Epipe service. ESM traffic from the host uses the subscriber queues that the host is configured to use, and the backhaul traffic uses the SAP queues. If **profile-traffic-only** is configured, then all traffic uses the SAP queues.

The operational status of the Epipe with a matching service ID can only be up if the corresponding Layer 3 SAP's operational status is up. The operational or administrative status of the Epipe does not affect the status of the Layer 3 service SAP. If the Layer 3 service SAP is up, but the backhaul Epipe is down, then the system continues to redirect packets to the Epipe, but they are dropped by the Epipe service. The status of the Epipe service is indicated in the output of the **show service service-using** command.

5.2 MLPPPoE, MLPPP(oE)oA with LFI on LNS

MLPPPoX is generally used to address bandwidth constraints in the last mile. The following are other uses for MLPPPoX:

- To increase bandwidth in the access network by bundling multiple links/VCs together. For example it is less expensive for a customer with an E1 access to add another E1 link in order to increase the access b/w, rather than to upgrade to the next circuit speed (E3).
- LFI on a single link to prioritize small packet size traffic over traffic with large size packets. This is needed in the upstream and downstream direction.

PPPoE and PPPoEoA/PPPoA v4/v6 host types are supported.

5.2.1 Terminology

The term MLPPPoX is used to reference MLPPP sessions over ATM transport (oA), Ethernet over ATM transport (oEoA) or Ethernet transport (oE). Although MLPPP in subscriber management context is not supported natively over PPP/HDLC links, the terms MLPPP and MLPPPoX terms can be used interchangeably. The reason for this is that link bundling, MLPPP encapsulation, fragmentation and interleaving can be in a broader scope observed independently of the transport in the first mile. However, MLPPPoX terminology prevails in this section in an effort to distinguish MLPPP functionality on an ASAP MDA (outside of ESM) and MLPPPoX in LNS (inside of ESM).

Terms speed and rate are interchangeably used throughout this section. Usually, speed refers to the speed of the link in general context (high or low) while rate quantitatively describes the link speed and associates it with the specific value in b/s.

5.2.2 LNS MLPPPoX

This functionality is supported through LNS on BB-ISA. LNS MLPPPoX can be used then as a workaround for PTA deployments, whereby LAC and LNS can be run back-to-back in the same system (connected via an external loop or a VSM2 module), and thus locally terminate PPP sessions.

MLPPPoX can:

- Increase bandwidth in the last mile by bundling multiple links together.

- LFI/reassembly over a single MLPPPoX capable link (plain PPP does not support LFI).

Typical MLPPPoA deployment is shown in Figure 35.

Figure 35 **Typical MLPPPoA Deployment**

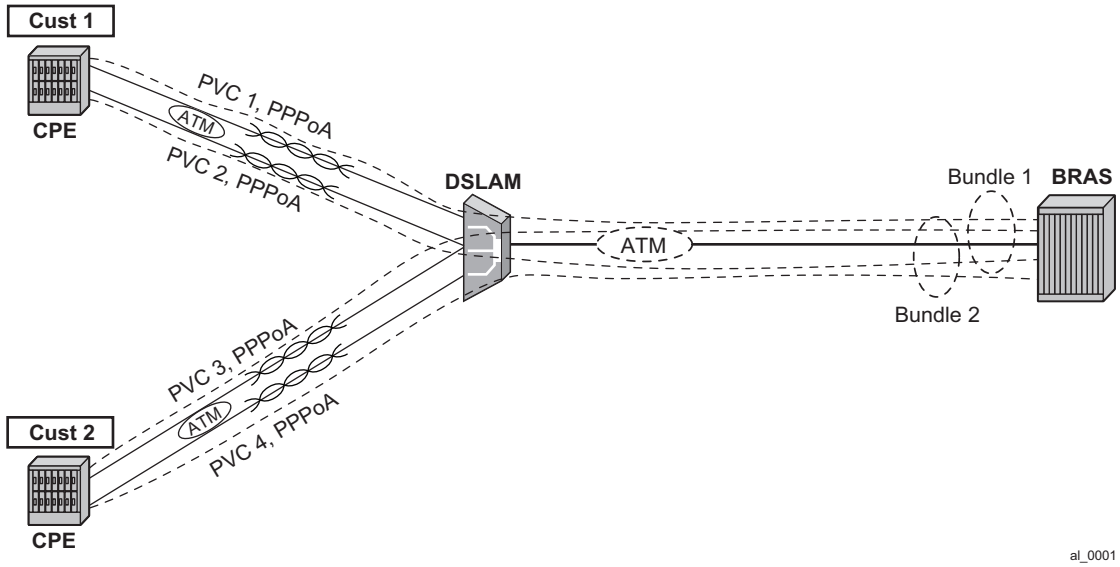

5.2.3 MLPPP Encapsulation

Once the MLPPP bundle is created in the 7750 SR, traffic can be transmitted by using MLPPP encapsulation. However, MLPPP encapsulation is not mandatory over an MLPPP bundle.

MLPPP header is primarily required for sequencing the fragments. But in case that a packet is not fragmented, it can be transmitted over the MLPPP bundle using either plain PPP encapsulation or MLPPP encapsulation. MLPPP encapsulation for fragmented traffic is shown in Figure 36.

Figure 36 **MLPPP Encapsulation**

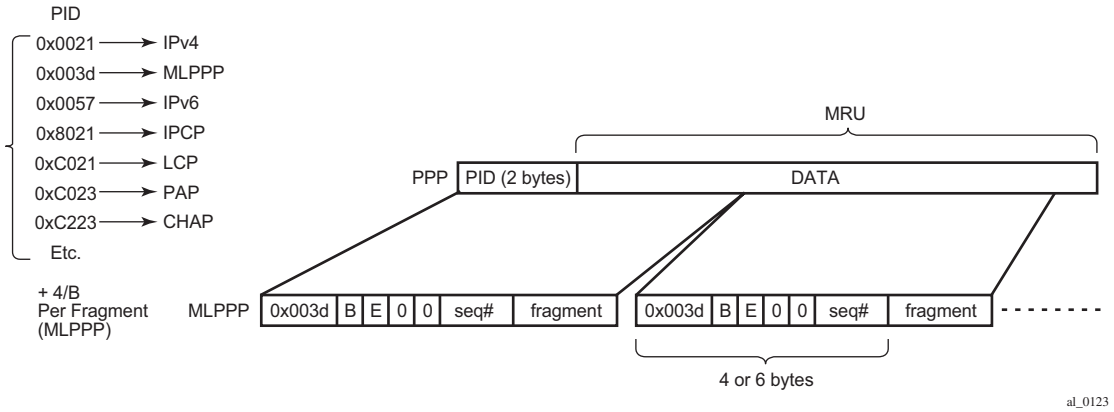

5.2.4 MLPPPoX Negotiation

MLPPPoX is negotiated during the LCP session negotiation phase by the presence of the Max-Received-Reconstructed Unit (MRRU) field in the LCP ConfReq. MRRU option is a mandatory field required in MLPPPoX negotiation. It represents the maximum number of octets in the Information field (Data part in Figure 36) of a reassembled packet. The MRRU value negotiated in the LCP phase must be the same on all member links and it can be greater or lesser than the PPP negotiated MRU value of each member link. This means that the reassembled payload of the PPP packet can be greater than the transmission size limit imposed by individual member links within the MLPPPoX bundle. Packets are always be fragmented so that the fragments are within the MRU size of each member link.

Another field that could be optionally present in an MLPPPoX LCP Conf Req is an Endpoint Discriminator (ED). Along with the authentication information, this field can be used to associate the link with the bundle.

The last MLPPPoX negotiated option is the Short Sequence Number Header Format Option which allows the sequence numbers in MLPPPoX encapsulated frames/ fragments to be 12-bit long (instead 24-bit long, by default).

Once the multilink capability is successfully negotiated via LCP, PPP sessions can be bundled together over MLPPPoX capable links.

The basic operational principles are:

- LCP session is negotiated on each physical link with MLPPPoX capabilities between the two nodes.

- Based on the ED and/or the authentication outcome, a bundle is created. A subsequent IPCP negotiation is conveyed over this bundle. User traffic is sent over the bundle.
- If a new link tries to join the bundle by sending a new MLPPPoX LCP Conf Request, the LCP session is negotiated, authentication performed and the link is placed under the bundle containing the links with the same ED and/or authentication outcome.
- IPCP/IPv6CP is in the whole process negotiated only once over the bundle. This negotiation occurs at the beginning, when the first link is established and MLPPPoX bundle created. IPCP and IPc6CP messages are transmitted from the 7750 SR LNS without MLPPPoX encapsulation, while they can be received as MLPPPoX encapsulated or non-MLPPPoX encapsulated.

5.2.5 Enabling MLPPPoX

The lowest granularity at which MLPPPoX can be enabled is an L2TP tunnel. An MLPPPoX enabled tunnel is not limited to carrying only MLPPPoX sessions but can carry normal PPP(oE) sessions as well.

In addition to enabling MLPPPoX on the session terminating node LNS, MLPPPoX can also be enabled on the LAC via PPP policy. The purpose of enabling MLPPPoX on the LAC is to negotiate MLPPPoX LCP parameters with the client. Once the LAC receives the MRRU option from the client in the initial LCP ConfReq, it changes its tunnel selection algorithm so that all sessions of an MLPPPoX bundle are mapped into the same tunnel.

The LAC negotiates MLPPPoX LCP parameters regardless of the transport technology connected to it (ATM or Ethernet). LCP negotiated parameters are passed by the LAC to the LNS via Proxy LCP in ICCN message. In this fashion the LNS has an option to accept the LCP parameters negotiated by the LAC or to reject them and restart the negotiation directly with the client.

The LAC transparently passes session traffic handed to it by the LNS in the downstream direction and the MLPPPoX client in the upstream direction. The LNS and the MLPPPoX client performs all data processing functions related to MLPPPoX such as fragmentation and interleaving.

Once the LCP negotiation is completed and the LCP transition into an open state (configuration ACKs are sent and received), the Authentication phase on the LAC begins. During the Authentication phase the L2TP parameters become known (l2tp group, tunnel, and so on), the session is extended by the LAC to the LNS via L2TP. If the Authentication phase does not return L2TP parameters, the session is terminated because the 7750 SR does not support directly terminated MLPPPoX sessions.

In the case that MLPPPoX is not enabled on the LAC, the LAC negotiates plain PPP session with the client. In case that the client accepts plain PPP instead of MLPPPoX as offered by the LAC, when the session is extended to the LNS, the LNS re-negotiates MLPPPoX LCP with the client on a MLPPPoX enabled tunnel. The LNS learns about the MLPPPoX capability of the client via Proxy LCP message in ICCN (first Conf Req received from the client is also send in Proxy LCP). If the there is no indication of the MLPPPoX capability of the client, the LNS establishes a plain PPP(oE) session with the client.

Note: There is no dependency between ATM autosensing on LAC and MLPPPoX because autosensing operates on a lower layer than PPP (LCP).

5.2.6 Link Fragmentation and Interleaving (LFI)

The purpose of LFI is to ensure that short high priority packets are not delayed by the transmission delay of large low priority packets on slow links.

For example it takes ~150ms to transmit a 5000B packet over a 256 kb/s link, while the same packet is transmitted in only 40us over a 1G link (~4000 times faster transmission). To avoid the delay of a high priority packet by waiting in the queue while the large packet is being transmitted, the large packet can be segmented into smaller chunks. The high priority packet can be then interleaved with the smaller fragments. This approach can significantly reduce the delay of high priority packets.

The interleaving functionality is only supported on MLPPPoX bundles with a single link. If more than one link is added into a interleaving capable MLPPPoX bundle, then interleaving is internally disabled and the tmnxMlpppBundleIndicatorsChange trap generated.

With interleaving enabled on an MLPPPoX enabled tunnel, the following session types are supported:

- Multiple LCP sessions tied into a single MLPPPoX bundle. This scenario assumes multiple physical links on the client side. Theoretically it would be possible to have multiple sessions running over the same physical link in the last mile. For example, two PPPoE sessions going over the same Ethernet link in the last mile, or two ATM VCs o2q1231qa23n the same last mile link. Whichever the case might be, the LAC/LNS is unaware of the physical topology in the last mile (single or multiple physical links). Interleaving functionality is internally disabled on such MLPPPoX bundle.
- A single LCP session (including dual stack) over the MLPPPoX bundle. This scenario assumes a single physical link on the client side. Interleaving is supported on such single session MLPPPoX bundle as long as the conditions for interleaving are met. Those conditions are governed by max-fragment-delay parameter and calculation of the fragment size as described in subsequent sections.
- An LCP session (including dual stack) over a plain PPP/PPPoE session. This type of session is a regular PPP(oE) session outside of any MLPPPoX bundle and therefore its traffic is not MLPPPoX encapsulated.

Packets on an MLPPPoX bundle are MLPPPoX encapsulated unless they are classified as high priority packets when interleaving is enabled.

5.2.6.1 MLPPPoX Fragmentation, MRRU and MRU Considerations

MLPPPoX in the 7750 SR is concerned with two MTUs:

- **bundle-mtu** determines the maximum length of the original IP packet that can be transmitted over the entire bundle (collection of links) before any MLPPPoX processing takes place on the transmitting side. This is also the maximum size of the IP packet that the receiving node can accept once it de-encapsulates and assembles received MLPPPoX fragments of the same packet. Bundle-mtu is relevant in the context of the collection of links.
- **link-mtu** determines the maximum length of the payload before it is PPP encapsulated and transmitted over an individual link within the bundle. Link-mtu is relevant in the context of the single link within the bundle.

Assuming that the CPE advertised MRRU and MRU values are smaller than any configurable mtu on MLPPPoX processing modules in 7750 SR (carrier IOM and BB-ISA), the bundle-mtu and the link-mtu are based on the received MRRU and MRU values, respectively. For example, the bundle-mtu is set to the received MRRU value while link-bundle is set to the MRU value minus the MLPPPoX encapsulation overhead (4 or 6 bytes).

In addition to mtu values, fragmentation requires a fragment length value for each MLPPP bundle on LNS. This fragment length value is internally calculated according to the following parameters;

- Minimum desired transmission delay in the last mile.
- Fragment “payload to encapsulation overhead” efficiency ratio.
- Various MTU sizes in the 7750 SR dictated mainly by received MRU, received MRRU and configured PPP MTU under the following hierarchy:
  - configure subscriber-mgmt ppp-policy ppp-mtu (ignored on LNS)
  - configure service vprn l2tp group ppp mtu
  - configure service vprn l2tp group tunnel ppp mtu
  - configure router l2tp group ppp mtu
  - configure router l2tp group tunnel ppp mtu

The decision whether to fragment and encapsulate a packet in MLPPPoX depends on the mode of operation, the packet length and the packet priority as follows:

LFI Case — When Interleave is enabled in a bundle, low priority packets are always MLPPPoX encapsulated. If a low-priority packet’s length exceeds the internally calculated Fragment Length, the packet is MLPPPoX fragmented and encapsulated. High priority packets whose length is smaller than the link-mtu is PPP encapsulated and transmitted without MLPPP encapsulation.

Non-LFI Case — When Interleave is disabled in a bundle, all packets are MLPPPoX encapsulated. If a packet’s length exceeds the internally calculated fragment length, the packet is MLPPPoX fragmented and encapsulated.

A packet of the size greater than the link-mtu cannot be natively transmitted over an MLPPPoX bundle. This packet is MLPPPoX encapsulated and consequently fragmented. This is regardless of the priority of the packet in interleaving case or whether the fragmentation is enabled or disabled.

In cases where MLPPPoX fragmentation is disabled with the no max-fragment-delay command, it is expected that packets are not MLPPPoX fragmented but rather only MLPPPoX encapsulated in order to be load balanced over multiple physical links in the last mile. However, even if MLPPPoX fragmentation is disabled, it is possible that fragmentation occurs under certain circumstances. This behavior is related to the calculation of the MTU values on an MLPPPoX bundle.

Consider an example where received MRRU value sent by CPE is 1500B while received MRU is 1492B. In this case, our bundle-mtu is set to 1500B and our link-mtu is set to 1488B (or 1486B) to allow for the additional 4/6B of MLPPPoX encapsulation overhead. Consequently, IP payload of 1500B can be transmitted over the bundle but only 1488B can be transmitted over any individual link. In case that an IP packet with the size between 1489B and 1500B needs to be transmitted from 7750 SR towards the CPE, this packet would be MLPPPoX fragmented in 7750 SR as dictated by the link-mtu. This is irrespective of whether MLPPPoX fragmentation is enabled or disabled (as set by no max-fragment-delay flag).

To entirely avoid MLPPPoX fragmentation in this case, the received MRRU sent by CPE should be lower than the received MRU for the length of the MLPPPoX header (4 or 6 bytes). In this case, for IP packets larger than 1488B, IP fragmentation would occur (assuming that DF flag in the IP header allows it) and MLPPPoX fragmentation would be avoided.

On the 7750 SR side, it is not possible to set different advertised MRRU and MRU values with the ppp-mtu command. Both MRRU and MRU advertised values adhere to the same configured ppp mtu value.

5.2.7 LFI Functionality Implemented in LNS

As mentioned in the previous section, LFI on LNS is implemented only on MLPPPoX bundles with a single LCP session.

There are two major tasks (Most of this is also applicable to non-lfi case. The only difference between lfi and non-lfi is that there is no artificial delay performed in non-lfi case) associated with LFI on the LNS:

- Executing subscriber QoS in the carrier IOM based on the last mile conditions. The subscriber QoS rates are the last mile on-the-wire rates. Once traffic is QoS conditioned, it is sent to the BB-ISA for further processing.
- Fragmentation and artificial delay (queuing) of the fragments so that high priority packets can be injected in-between low priority fragments (interleaved). This operation is performed by the BB-ISA.

Examine an example to further clarify functionality of LFI. The parameters, conditions and requirements that are used in our example to describe the desired behavior are the following:

- High priority packets must not be delayed for more than 50ms in the last mile due to the transmission delay of the large low priority packets. Considering that tolerated end-to-end VoIP delay must be under 150ms, limiting the transmission delay to 50ms on the last mile link is a reasonable choosing.

- The link between the LNS and LAC is 1Gb/s Ethernet.
- The last mile link rate is 256 kb/s.
- Three packets arrive back-to-back on the network side of the LNS (in the downstream direction). The large 5000B low priority packet P1 arrives first, followed by two smaller high priority packets P2 and P3, each 100B.

Note: Packets P1, P2 and P3 can be originated by independent sources (PCs, servers, etc.) and therefore can theoretically arrive in the LNS from the network side back-to-back at the full network link rate (10Gb/s or 100Gb/s).

- The transmission time on the internal 10G link between the BB-ISA and the carrier IOM for the large packet (5000B) is 4us while the transmission time for the small packet (100B) is 80ns.
- The transmission time on the 1G link (LNS->LAC) for the large packet (5000B) is 40us while the transmission time for the small packet (100B) is 0.8us.
- The transmission time in the last mile (256 kb/s) for the large packet is ~150ms while the transmission time for the small packet on the same link is ~3ms.
- Last mile transport is ATM.

To satisfy the delay requirement for the high priority packets, the large packets are fragmented into three smaller fragments. The fragments are carefully sized so that their individual transmission time in the last mile does not exceed 50ms. After the first 50ms interval, there is an opportunity to interleave the two smaller high priority packets.

This entire process is further clarified by the five points (1-5) in the packet route from the LNS to the Residential Gateway (RG) as depicted in Figure 37.

The five points are:

1. Last Mile QoS Awareness in the LNS
2. BB-ISA Processing
3. LNS-LAC Link
4. AN-RG Link
5. Home Link

Figure 37 **Packet Route from the LNS to the RG**

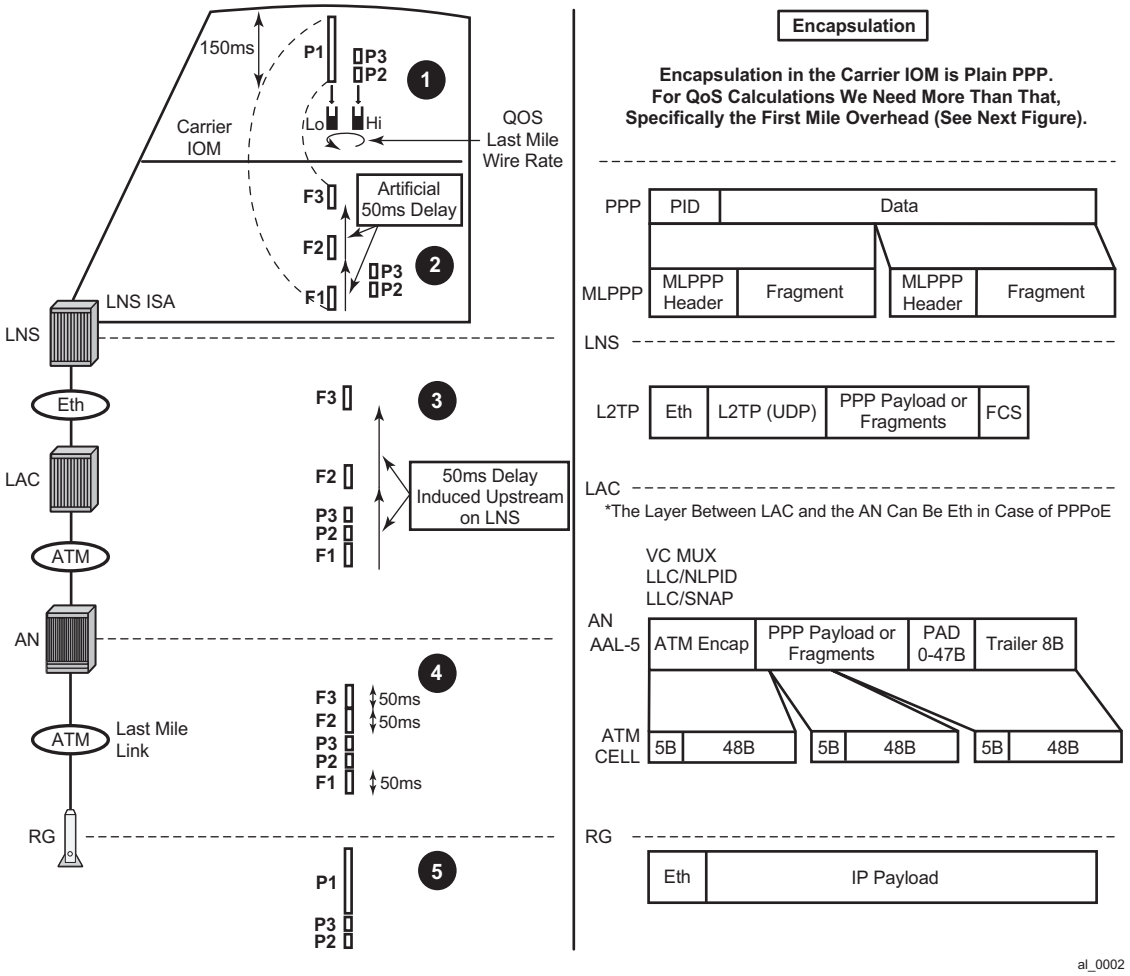

5.2.7.1 Last Mile QoS Awareness in the LNS

By implementing MLPPPoX in LNS, we are effectively transferring the traffic treatment functions (QoS/LFI) of the last mile to the node (LNS) that is multiple hops away.

The success of this operation depends on the accuracy at which we can simulate the last mile conditions in the LNS. The assumption is that the LNS is aware of the two most important parameters of the last mile:

- The last mile encapsulation — This is needed for the accurate calculation of the overhead associated of the transport medium in the last mile for traffic shaping and interleaving.

- The last mile link rate — This is crucial for the creation of artificial congestion and packet delay in the LNS.

The subscriber QoS in the LNS is implemented in the carrier IOM and is performed on a per packets basis before the packet is handed over to the BB-ISA. Per packet, rather than per fragment QoS processing ensures a more efficient utilization of network resources in the downstream direction. Discarding fragments in the LNS would have detrimental effects in the RG as the RG would be unable to reconstruct a packet without all of its fragments.

High priority traffic within the bundle is classified into the high priority queue. This type of traffic is not MLPPPoX encapsulated unless its packet size exceeds the link MTU as described in MLPPPoX Fragmentation, MRRU and MRU Considerations. Low priority traffic is classified into a low priority queue and is always MLPPPoX encapsulated. In case that the high priority traffic becomes MLPPPoX encapsulated or fragmented, the MLPPPoX processing module (BB-ISA) considers it as low-priority. The assumption is that the high priority traffic is small in size and consequently MLPPPoX encapsulation or fragmentation an degradation in priority can be avoided. The aggregate rate of the MLPPPoX bundle is on-the-wire rate of the last mile as shown in Figure 38.

ATM on-the-wire overhead for non-MLPPPoX encapsulated high priority traffic includes:

- ATM encapsulation (VC-MUX, LLC/NLPID, LLC/SNAP).
- AAL5 trailer (8B).
- AAL5 padding to 48B cell boundary (this makes the overhead dependent on the packet size).
- Multiplication by 53/48 to account for the ATM cell headers.

For low priority traffic which is always MLPPPoX encapsulated, an additional overhead related to MLPPPoX encapsulation and possibly fragmentation must be added. In other words, each fragment carries ATM+MLPPPoX overhead.

Note: Avoid the 48B boundary padding for all fragments except the last one. This can be done by choosing the fragment length so that it is aligned on the 48B boundary (rounded down if based on max-fragment-delay or rounded up if based on the encapsulation/ utilization.

Figure 38 **Last Mile Encapsulation**

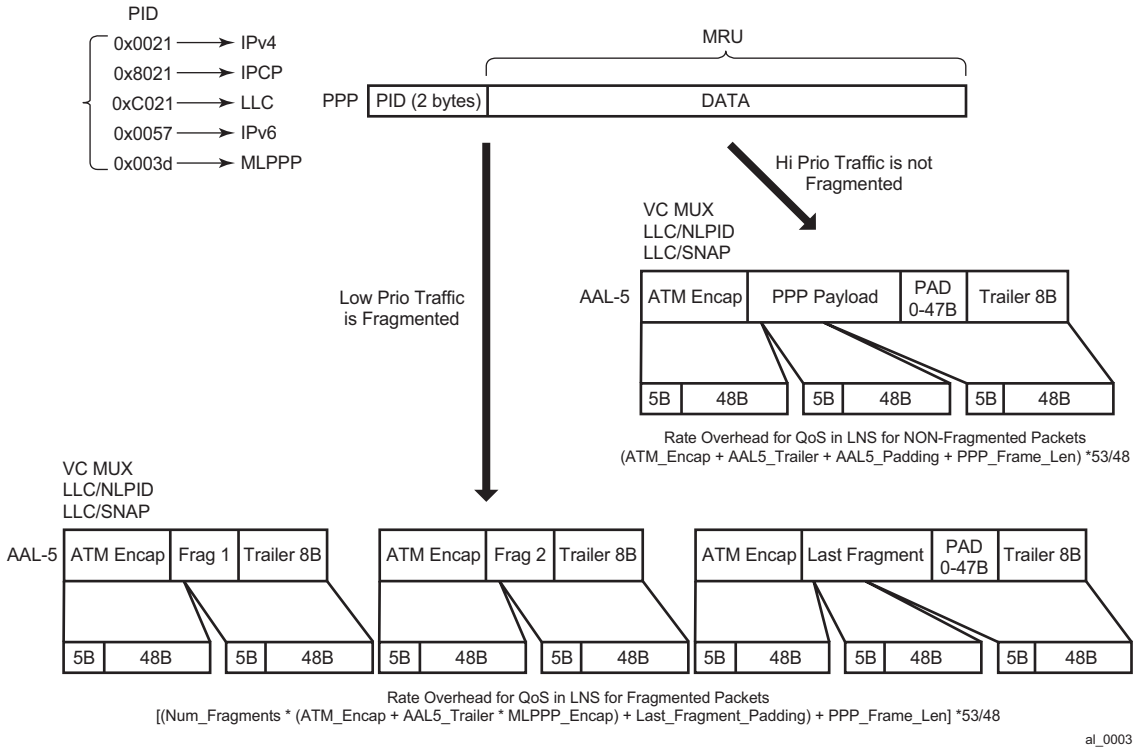

For Ethernet in the last mile, our implementation always assures that the fragment size plus the encapsulation overhead is always larger or equal to the minimum Ethernet packet length (64B).

5.2.7.2 BB-ISA Processing

MLPPPoX encapsulation, fragmentation and interleaving are performed by the LNS in BB-ISA. If we refer to our example, a large low priority packet (P1) is received by the BB-ISA, immediately followed by the two small high priority packets (P2 and P3). Since our requirement stipulates that there is no more than 50ms of transmission delay in the last mile (including on-the-wire overhead), the large packet must be fragmented into three smaller fragments each of which will not cause more than 50ms of transmission delay.

The BB-ISA would normally send packets/fragments to the carrier IOM at the rate of 10Gb/s. In other words, by default the three fragments of the low priority packet would be sent out of the BB-ISA back-to-back at the very high rate before the high priority packets even arrive in the BB-ISA. In order to interleave, the BB-ISA must simulate the last mile conditions by delaying the transmission of the fragments. The fragments are be paced out of the BB-ISA (and out of the box) at the rate of the last mile. High priority packets can be injected in front of the fragments while the fragments are being delayed.

In Figure 37 (point 2) the first fragment F1 is sent out immediately (transmission delay at 10G is in the 1us range). The transmission of the next fragment F2 is delayed by 50ms. While the transmission of the second fragment F2 is being delayed, the two high priority packets (P1 and P2 in red) are received by the BB-ISA and are immediately transmitted ahead of fragments F2 and F3. This approach relies on the imperfection of the IOM shaper which is releasing traffic in bursts (P2 and P3 right after P1). The burst size is dependent on the depth of the rate token bucket associated with the IOM shaper.

Note: By the time the second fragment F2 is transmitted, the first fragment F1 has traveled a long way (50ms) on high rate links towards the Access Node (assuming that there is no queuing delay along the way), and its transmission on the last mile link has already begun (if not already completed).

This is not applicable for this discussion, but nonetheless worth noticing is that the LNS BB-ISA also adds the L2TP encapsulation to each packet/fragment. The L2TP encapsulation is removed in the LAC before the packet/fragment is transmitted towards the AN.

5.2.7.3 LNS-LAC Link

This is the high rate link (1Gb/s) on which the first fragment F1 and the two consecutive high priority packets, P2 and P3, are sent back-to-back by the BB-ISA

(BB-ISA->carrier IOM->egress IOM-> out-of-the-LNS).

The remaining fragments (F2 and F3) are still waiting in the BB-ISA to be transmitted. They are artificially delayed by 50ms each.

Additional QoS based on the L2TP header can be performed on the egress port in the LNS towards the LAC. This QoS is based on the classification fields inside of the packet/fragment headers (DSCP, dot1.p, EXP).

Note: The LAC-AN link is not really relevant for the operation of LFI on the LNS. This link can be either Ethernet (in case of PPPoE) or ATM (PPPoE or PPP). The rate of the link between the LAC and the AN is still considered a high speed link compared to the slow last mile link.

5.2.7.4 AN-RG Link

Finally, this is the slow link of the last mile, the reason why LFI is performed in the first place. Assuming that LFI played its role in the network as designed, by the time the transmission of one fragment on this link is completed, the next fragment arrives just in time for unblocked transmission. In between the two fragments, we can have one or more small high priority packets waiting in the queue for the transmission to complete.

Note: On the AN-RG link in Figure 37 that packets P2 and P3 are ahead of fragments F2 and F3. Therefore the delay incurred on this link by the low priority packets is never greater than the transmission delay of the first fragment (50ms). The remaining two fragments, F2 and F3, can be queued and further delayed by the transmission time of packets P2 and P3 (which is normally small, in our example 3ms for each).

Note: If many low priority packets are waiting in the queue, then they would have caused delay and would have further delayed the fragments that are in transit from the LNS to the LAC. This condition is normally caused by bursts and it should clear itself out over time.

5.2.7.5 Home Link

High priority packets P2 and P3 are transmitted by the RG into the home network ahead of the packet P1 although the fragment F1 has arrived in the RG first. The reason for this is that the RG must wait for the fragments F2 and F3 before it can re-assemble packet P1.

5.2.7.6 Optimum Fragment Size Calculation by LNS

Fragmentation in LFI is based on the optimal fragment size. LNS implementation calculates the two optimal fragment sizes, based on two different criteria:

- Optimal fragment size based on the payload efficiency of the fragment given the fragmentation/transportation header overhead associated with the fragment encapsulation based fragment size.
- Optimal fragment size based on the maximum transmission delay of the fragment set by configuration delay based fragment size.

At the end, only one optimal fragment size is selected. The actual fragment's length is the optimal fragment size.

- The parameters required to calculate the optimal fragment sizes are known to the LNS either via configuration or via signaling. These, in-advance known parameters are:
- Last mile maximum transmission delay (max-fragment-delay obtained via CLI)
- Last mile ATM Encapsulation (in our example the last mile is ATM but in general it can be Ethernet for MLPPPoE)
- MLPPP encapsulation length (depending on the fragment sequence number format)
- The last mile on-the-wire rate for the MLPPPoX bundle

Examine closer each of the two optimal fragment sizes.

5.2.7.6.1 Encapsulation Based Fragment Size

One needs to be mindful of the fact that fragmentation may cause low link utilization. In other words, during fragmentation a node may end up transporting mainly overhead bytes in the fragment as opposed to payload bytes. This would only intensify the problem that fragmentation is intended to solve, especially on an ATM access link that tend to carry larger encapsulation overhead.

To reduce the overhead associated with fragmentation, the following is enforced in the 7750 SR:

The minimum fragment payload size is at least 10times greater than the overhead (MLPPP header, ATM Encapsulation and AAL5 trailer) associated with the fragment.

The optimal fragment length (including the MLPPP header, the ATM Encapsulation and the AAL5 trailer) is a multiple of 48B. Otherwise, the AAL5 layer would add an additional 48B boundary padding to each fragment which would unnecessary expand the overhead associated with fragmentation. By aligning all-but-last fragments to a 48B boundary, only the last fragment potentially contains the AAL5 48B boundary padding which is no different from a non-fragmented packet. All fragments, except for the last fragment, are referred to as non-padded fragments. The last fragment is padded if it is not already natively aligned to a 48B boundary.

As an example, calculate the optimal fragment size based on the encapsulation criteria with the maximum fragment overhead of 22B. To achieve >10x transmission efficiency the fragment payload size must be 220B (10*22B). To avoid the AAL5 padding, the entire fragment (overhead + payload) is rounded UP on a 48B boundary. The final fragment size is 288B [22B + 22B*10 + 48B_allignment].

In conclusion, an optimal fragment size was selected that carries the payload with at least 90% efficiency. The last fragment of the packet cannot be artificially aligned on a 48B boundary (it is a natural reminder), so it is be padded by the AAL5 layer. Therefore, the efficiency of the last fragment is be less than 90% in our example. In the extreme case, the efficiency of this last fragment may be only 2%.

Note: The fragment size chosen in this manner is purely chosen based on the overhead length. The maximum transmission delay did not play any role in the calculations.

For Ethernet based last mile, the CPM always makes sure that the fragment size plus encapsulation overhead is larger or equal to the minimum Ethernet packet length of 64B.

5.2.7.6.2 Fragment Size Based on the Max Transmission Delay

The first criterion in selecting the optimal fragment size based on the maximum transmission delay mandates that the transmission time for the fragment, including all overheads (MLPPP header, ATM encapsulation header, AAL5 overhead and ATM cell overhead) must be less than the configured max-fragment-delay time.

The second criterion mandates that each fragment, including the MLPPP header, the ATM Encapsulation header, the AAL5 trailer and the ATM cellification overhead be a multiple of 48B. The fragment size is rounded down to the nearest 48B boundary during the calculations in order to minimize the transmission delay. Aligning the fragment on the 48B boundary eliminates the AAL5 padding and therefore reduces the overhead associated with the fragment. The overhead reduction improves the transmission time and also increases the efficiency of the fragment.

Given these two criteria along with the configuration parameters (ATM Encapsulation, MLPPP header length, max-fragment-delay time, rate in the last mile), the implementation calculates the optimal non-padded fragment length as well as the transmission time for this optimal fragment length.

5.2.7.6.3 Selection of the Optimum Fragment Length

So far the implementation has calculated the two optimum fragment lengths, one based on the length of the MLPPP/transport encapsulation overhead of the fragment, the other one based on the maximum transmission delay of the fragment. Both of them are aligned on a 48B boundary. The larger of the two is chosen and the BB-ISA performs LFI based on this selected optimal fragment length.

5.2.8 Upstream Traffic Considerations

Fragmentation and interleaving is implemented on the originating end of the traffic. In other words, in the upstream direction the CPE (or RG) is fragmenting and interleaving traffic. There is no interleaving or fragmentation processing in the upstream direction in the 7750 SR. The 7750 SR are on the receiving end and is only concerned with the reassembly of the fragments arriving from the CPE. Fragments are buffered until the packet can be reconstructed. If all fragments of a packet are not received within a preconfigured timeframe, the received fragments of the partial packet are discarded (a packet cannot be reconstructed without all of its fragments). This time-out and discard is necessary in order to prevent buffer starvation in the BB-ISA. Two values for the time-out can be configured: 100ms and 1s.

5.2.9 Multiple Links MLPPPoX With No Interleaving

Interleaving over MLPPPoX bundles with multiple links are not supported. However, fragmentation is supported.

In order to preserve packet order, all packets on an MLPPPoX bundle with multiple links are MLPPPoX encapsulated (monotonically increased sequence numbers).

Multiclass MLPPP (RFC 2686, *The Multi-Class Extension to Multi-Link PPP*) is not supported.

5.2.10 MLPPPoX Session Support

The following session types in the last mile are supported:

- MLPPPoE — Single physical link or multilink. The last mile encapsulation is Ethernet over copper (This could be Ethernet over VDSL or HSDSL). The access rates (especially upstream) are still limited by the xDSL distance limitation and as such interleaving is required on a slow speed single link in the last mile. It is possible that the last mile encapsulation is Ethernet over fiber (FTTH) but in this case, users would not be concerned with the link speed to the point where interleaving and link aggregation is required.

Finally, this is the slow link of the last mile, the reason why LFI is performed in the first place. Assuming that LFI played its role in the network as designed, by the time the transmission of one fragment on this link is completed, the next fragment arrives just in time for unblocked transmission. In between the two fragments, we can have one or more small high priority packets waiting in the queue for the transmission to complete.

We can see on the AN-RG link in Figure 2 that packets P2 and P3 are ahead of fragments F2 and F3. Therefore the delay incurred on this link by the low priority packets is never greater than the transmission delay of the first fragment (50ms). The remaining two fragments, F2 and F3, can be queued and further delayed by the transmission time of packets P2 and P3 (which is normally small, in our example 3ms for each).

Note: If many low priority packets were waiting in the queue, then they would have caused delay for each other and would have further delayed the fragments in transit from the LNS to the LAC. This condition is normally caused by bursts and it should clear itself out over time.

Figure 39 **MLPPPoE — Multiple Physical Links**

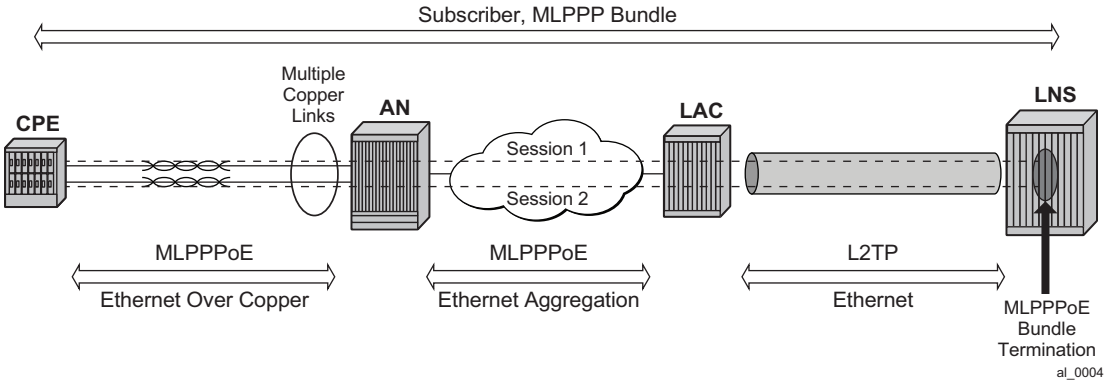

Figure 40 **MLPPPoE — Single Physical Link**

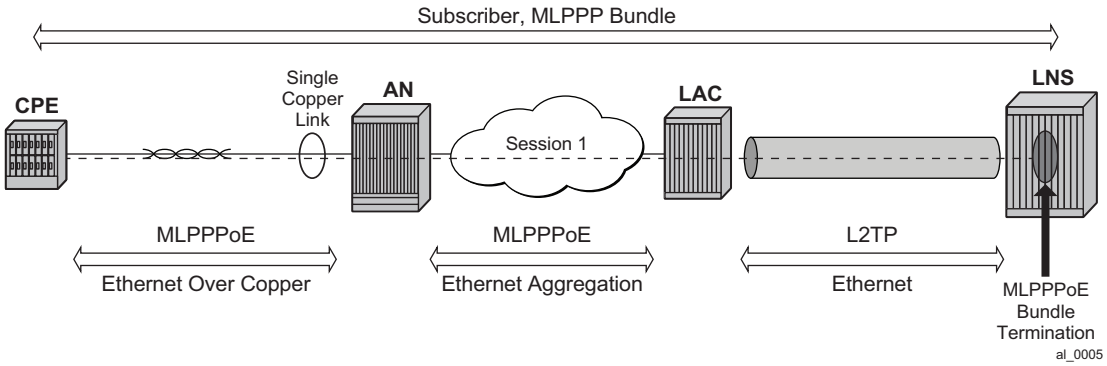

- MLPPP(oEo)A — A single physical link or multilink. The last mile encapsulation is ATM over xDSL.

Figure 41 **MLPPP(oE)oA — Multiple Physical Links**

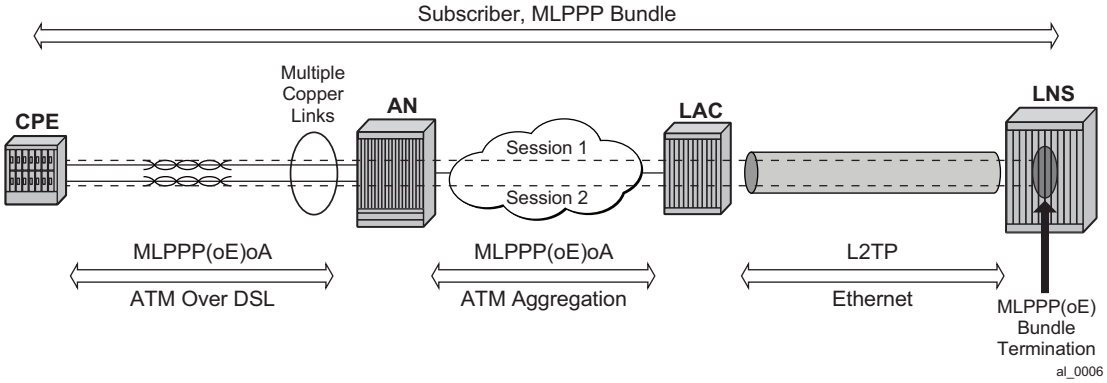

Figure 42 **MLPPP(oE)oA — Single Physical Link**

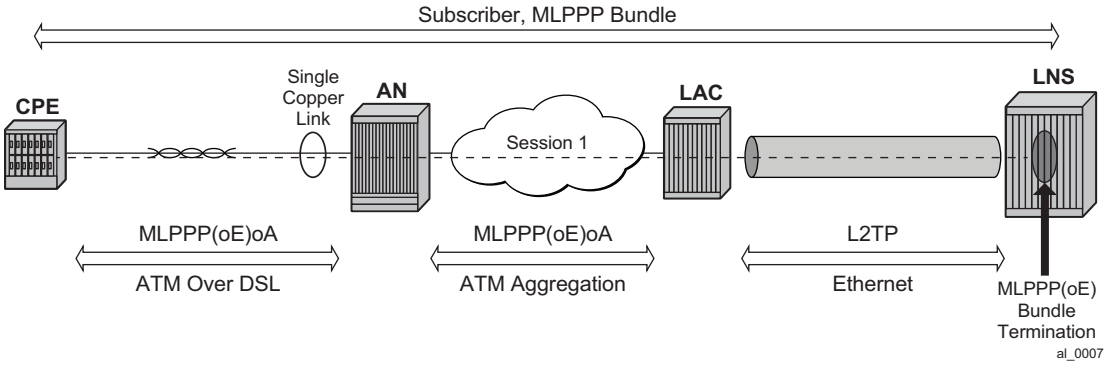

Some other combinations are also possible (ATM in the last mile, Ethernet in the aggregation) but they all come down to one of the above models that are characterized by:

- Ethernet or ATM in the last mile.
- Ethernet or ATM access on the LAC.
- MLPPP/PPPoE termination on the LNS

5.2.11 Session Load Balancing Across Multiple BB-ISAs

PPP/PPPoE sessions are by default load balanced across multiple BB-ISAs (max 6) in the same group. The load balancing algorithm considers the number of active session on each BB-ISA in the same group. The load balancing algorithm does not take into account the number of queues consumed on the carrier IOM. Therefore a session can be refused if queues are depleted on the carrier IOM even though the BB-ISA may be lightly loaded in terms of the number of sessions that is hosting.

With MLPPPoX, it is important that multiple sessions per bundle be terminated on the same LNS BB-ISA. This can be achieved by per tunnel load balancing mode where all sessions of a tunnel are terminated in the same BB-ISA. Per tunnel load balancing mode is mandatory on LNS BB-ISAs that are in the group that supports MLPPPoX.

On the LAC side, all sessions in an MLPPPoX bundle are automatically assigned to the same tunnel. In other words an MLPPPoX bundle is assigned to the tunnel. There can be multiple tunnels created between the same pair of LAC/LNS nodes.

5.2.12 BB-ISA Hashing Considerations

All downstream traffic on an MLPPPoX bundle with multiple links is always MLPPPoX encapsulated. Some traffic is fragmented and served in a octet oriented round robin fashion over multiple member links. However, fragments are never delayed in case that the bundle contains multiple links.

In a per fragment/packet load sharing algorithm, there is always the possibility that there is uneven load utilization between the member links. A single link overload can go unnoticed in the network all the way to the Access Node. The access node is the only node in the network that actually has multiple physical links connected to it. All other session-aware nodes (LAC and LNS) only see MLPPPoX as a bundle with multiple sessions without any mechanism to shape traffic per physical link. Other nodes in this case being 7750 SRs. Other vendors may have the ability to condition (shape) traffic per session.

If one of the member sessions is perpetually overloaded by the LNS, traffic is dropped in the last mile since the corresponding physical link cannot absorb traffic beyond its physical capabilities. This would have detrimental effects on the whole operation of the MLPPPoX bundle. To prevent this perpetual overloading of the member links that can be caused by per packet/fragment load balancing scheme, the load balancing scheme that takes into account the number of octets transmitted over each member link. The octet counter of a new link is initialized to the lowest value of any existing link counter. Otherwise the load balancing mechanism would show significant bias towards the new link until the byte counter catches up with the rest of the links.

5.2.13 Last Mile Rate and Encapsulation Parameters

The last mile rate information along with the encapsulation information is used for fragmentation (to determine the maximum fragment length) and interleaving (delaying fragments in the BB-ISA). In addition, the aggregate subscriber rate (aggregate-rate-limit) on the LNS is automatically adjusted based on the last mile link rate and the number of links in the MLPPPoX bundle.

Downstream Data Rate in the Last Mile

The subscriber aggregate rates (agg-rate-limit) used in (H)QoS on the carrier IOM and in the BB-ISA (for interleaving) must be wire based in the last mile. This rule applies equally to both, the LAC and LNS.

The last mile on-the-wire rates of the subscriber can be submitted to the LAC and the LNS via various means. The following discusses the break down on how the last mile wire rates are passed to each entity:

LAC

The last mile link rate is taken via the following methods in the order of listed priority:

- LUDB — rate-down command under the host hierarchy in LUDB.
- RADIUS Alc-Access-Loop-Rate-Down VSA. Although this VSA is stored in the state of plain PPP(oE) sessions (MLPPPoX bundled or not), it is applicable only to MLPPPoX bundles.
- PPPoE tags — Vendor Specific Tags (RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*; tag type 0x0105; tag value is Enterprise Number 3561 followed by the TLV sub-options as specified in TR-101 -> Actual Data Rate Downstream 0x82)

As long as the link rate information is available in the LAC, it is always passed to the LNS in the ICRQ message using the standard L2TP encoding. This cannot be disabled.

In addition, an option is available to control the source of the rate information can be conveyed to the LNS via TX Connect Speed AVP in the ICCN message. This can be used for compatibility reasons with other vendors that can only use TX Connect Speed to pass the link rate information to the LNS. By default, the maximum port speed (or the sum of the maximum speeds of all member ports in the LAG) is reported in TX Connect Speed. Unlike the rate conveyed in ICRQ message, The TX Connect Speed content is configurable via the following command:

```
config>subscr-mgmt
    sla-profile <name>
        egress
            report-rate agg-rate-limit | scheduler <scheduler-name> | pppoe-actual-
rate | rfc5515-actual-rate
```

The report-rate configuration option dictates which rate is reported in the TX Connect Speed as follows:

- agg-rate-limit — statically configured agg-rate-limit value or RADIUS QoS override is reported
- scheduler <scheduler-name> — virtual schedulers are not supported in MLPPPoX
- pppoe-actual-rate — rate taken from PPPoE Tags is reported.

Note: Rate reported via RFC5515 can still be different if the source for both methods is not the same.

- rfc5515-actual-speed => the rate is taken from RFC5515.

The RFC 5515 relies on the same encoding as PPPoE tags (vendor id is ADSL Forum and the type for Actual Data Rate Downstream is 0x82).

Note: The two methods of passing the line rate to the LNS are using different message types (ICRQ and ICCN).

The LAC on the 7750 SR is not aware of MLPPPoX bundles. As such, the aggregate subscriber bandwidth on the LAC is configured statically via usual means (sub-profile, scheduler-policy) or dynamically modified via RADIUS. The aggregate subscriber (or MLPPPoX bundle) bandwidth on the LAC is not automatically adjusted according to the rates of the individual links in the bundle and the number of the links

in the bundle. As such, an operator must ensure that the statically provided rate value for aggregate-rate-limit is the sum of the bandwidth of each member link in the MLPPPoX bundle. The number of member links and their bandwidth must be therefore known in advance. The alternative is to have the aggregate rate of the MLPPPoX bundle set to a high value and rely on the QoS treatment performed on the LNS.

LNS

The sources of information for the last mile link rate on the LNS is taken in the following order:

- LUDB (during user authentication phase, same as in LAC)
- RADIUS (same as in LAC)
- ICRQ message — Actual Data Downstream Rate (RFC 5515)
- ICCN message — TX Connect Speed

There is no configuration option to determine the priority of the source of information for the last mile link rate. TX Connect Speed in ICCN message is only be taken into consideration as a last resort in absence of any other source of last mile rate information.

Once the last mile rate information is obtained, the subscriber aggregate rate aggregate-rate-limit is automatically adjusted to the minimum value of:

- The smallest link speed in the MLPPPoX bundle multiplied by the number of links in the bundle.
- Statically configured aggregate-rate-limit

The link speed of each link in the bundle must be the same, meaning, different link speeds within the bundle are not supported. When different link are received, speed values for last mile links within the bundle, the minimum received speed is adopted and apply it to all links.

When the obtained rate information from the last mile for a session within the MLPPP bundle is out of bounds (1 kb/s to 100 Mb/s), the session within the bundle is terminated.

Encapsulation

Wire-rates are dependent on the encapsulation of the link to which they apply. The last mile encapsulation information can be extracted via various means.

LAC

- Static configuration via LUDB.

- RADIUS — Alc-Access_Loop-Encap-Offset VSA.
- PPPoE tags — Vendor Specific Tags (RFC 2516; tag type 0x0105; tag value is Enterprise Number 3561 followed by the TLV sub-options as specified in TR-101 -> Actual Data Rate Downstream 0x82).

The LAC passes the line encapsulation information to the LNS via ICRQ message using the encoding defined in the RFC 5515.

LNS

The LNS extracts the encapsulation information in the following order:

- Static configuration via LUDB.
- RADIUS — Alc-Access-Loop-Encap-Offset VSA.
- ICRQ message (RFC 5515)

When the encapsulation information is not provided by any of the existing means (LUDB, RADIUS, AVP signaling, PPPoE Tags), then by default pppoa-null encapsulation is in effect. This applies to LAC and LNS.

5.2.14 Link Failure Detection

The link failure in the last mile is detected via the expiration of session keepalives (LCP). The LNS tears down the session over the failed link and notify the LAC via a CDN message.

5.2.15 CoA Support

CoA request for the subscriber aggregate-rate-limit change is honored on the LAC and the LNS.

CoA for the rate change of an individual link within the bundle is supported through the same VSA that can be used to initially assign the rate parameter to each member link. This is supported only on LNS. The rate override via CoA is applied to all active link members within the bundle.

Change of the access link parameters via CoA is be supported in the following fashion:

- Change of access loop encap: refused (NAK)
- Change of access loop rate down:

- On L2TP LAC session: refused (NAK). On LAC the access loop rate down is not locally used for any rate limiting function but instead it is just passed to the LNS at the beginning when the session is first established. Mid-session changes on LAC via CoA are not propagated to the LNS.
- On L2TP LNS session:
  - Plain session: ignored. The rate is stored in the MIB table but no rate limiting action is taken. In other words, this parameter is internally excluded from rate calculations and advertisements. However, it is shown in the output of the relevant show commands.
- Bundle session: applied on all link sessions. The aggregate rate limit of the bundle is set to the minimum of the:
- CoA obtained local loop down rate multiplied by the number of links in the bundle
- The aggregate rate limit configured statically or obtained via CoA.
- Fragment length is affected by this change. In case that interleaving is enabled on a single link bundle, the interleave interval is affected.
- Non-L2TP: ignored. The rate is stored in the MIB table but no rate limiting action is taken. In other words, this parameter is internally excluded from rate calculations and advertisements. However, it is shown in the output of the relevant show commands.

Similar behavior is exhibited if at mid session, the parameters are changed via LUDB with the exception of the rate-down parameter in LAC. If this parameter is changed on the LAC, all sessions are disconnected.

5.2.16 Accounting

Accounting counters on the LNS include all packet overhead (wire overhead from the last mile). There is only one accounting session per bundle.

On the LAC, there is one accounting session per pppoe session (link).

In tunnel-accounting mode there is one accounting session per link.

On LNS only the stop-link of the last link of the bundle carries all accounting data for the bundle.

5.2.17 Filters and Mirroring

Filters and mirrors (LI) are not supported on an MLPPPoX bundle on LAC. However, filters and ip-only mirror type are supported on the LNS.

5.2.18 PTA Considerations

Locally terminated MLPPPoX (PTA) solution is offered based on the LAC and the LNS hosted in the same system. An external loop (or VSM2) is used to connect the LAC to the LNS within the same box. The subscribers are terminated on the LNS.

5.2.19 QoS Considerations

5.2.19.1 Dual-Pass

HQoS and LFI are performed in two stages that involve double traversal (dual-pass) of traffic through the carrier IOM and the BB-ISA. The following are the functions performed in each pass:

- In the first pass through the carrier IOM, traffic is marked (dot1p bits) as high or low priority. This plays a crucial role in the execution of LFI in the BB-ISA.
- In the first pass through the BB-ISA this prioritization from the first step, is an indication (along with the internally calculated fragment size) of whether the traffic is interleaved (non MLPPP encapsulated) or not (MLPPP encapsulated). Consequently, the BB-ISA adds the necessary padding related to last mile wire overhead to each packet. This padding is in the second pass on the carrier IOM and performs last mile wire based QoS functions.
- In the second pass through the carrier IOM, the last-mile wire-based HQoS is performed based on the padding added in the first pass through the BB-ISA.
- In the second pass through the BB-ISA, the previously added overhead is stripped off and LFI/MLPPP encapsulation functions are performed.

5.2.19.2 Traffic Prioritization in LFI

The delivery of high priority traffic within predefined delay bounds on a slow speed last mile link is ensured by proper QoS classification and prioritization. High priority traffic is interleaved with low priority fragments on a single link MLPPPoX bundle with LFI enabled. The classification of traffic into proper (high or low priority) forwarding class is performed on the downstream ingress interface. However, traffic can be re-classified (re-mapped into another forwarding class) on the egress access interface of the carrier IOM, just before packets are transmitted to the BB-ISA for MLPPPoX processing. This can be achieved via QoS sap-egress policy referenced in the LNS sla-profile.

The priority of the forwarding class in regular QoS (on IOM) is determined by the properties of the queue to which the forwarding class is mapped. Expedited, non-expedited queue type, CIR and PIR rates. In contracts, traffic prioritization in LFI domain (in BB-ISA) is determined by the outer dot1p bits that are set by the carrier IOM while transmitting packets towards the BB-ISA. The outer dot1p bits are marked based on the forwarding class information determined by classification/re-classification on ingress/carrier IOM. This marking of outer dot1p bits in the Ethernet header between the carrier IOM and the BB-ISA is fixed and defined in the default sap-egress LNS ESM policy 65537. The marking definition is as follows:

```
FC be -> dot1p 0
FC l2 -> dot1p 1
FC af -> dot1p 2
FC l1 -> dot1p 3
FC h2 -> dot1p 4
FC ef -> dot1p 5
FC h1 -> dot1p 6
FC nc -> dot1p 7
```

In LFI (on BB-ISA), dot1p bits [0,1,2 and 3] are considered low priority while dot1p bits (4,5,6 and 7) are considered high priority. Consequently, forwarding classes BE, L2, AF and L1 are considered low priority while forwarding classes H2, EF, H1 and NC are considered high priority. High priority traffic is interleaved with low priority traffic. Assuming that the packet size does not exceed maximum fragment size.

The following describes the reference points in traffic prioritization for the purpose of LFI in the 7750 SR:

- Classification on downstream ingress interface (entrance point into the 7750 SR) - packets can be classified into one of the following eight forwarding classes: be, l2, af, l1, h2, ef, h1 and nc. Depending on the type of the ingress interface (access or network), traffic can be classified based on dot1p, exp, DSCP, TOS bits or ip-match criteria (dscp, dst-ip, dst-port, fragment, src-ip, src-port and protocol-id).

- Re-classification on downstream access egress interface between the carrier IOM and the BB-ISA - in the carrier IOM, downstream traffic can be re-classified into another forwarding class, just before it is forwarded to the BB-ISA. Re-classification on access egress is based on the same fields as on ingress except for the dot1p and exp bits since Ethernet or MPLS headers from ingress are not carried from ingress to egress.
- Marking on downstream access egress interface between the carrier IOM and the BB-ISA - once the forwarding class is available on the carrier IOM in the egress direction (towards BB-ISA), it is used to mark outer dot1p bits in the new Ethernet header that are used to transport the frame from the carrier IOM to the BB-ISA. The marking of the dot1p bits on the egress SAP between the carrier IOM and the BB-ISA cannot be changed for MLPPPoX even if the no qos-marking-from-sap command is configured under the sla-profile on egress.

5.2.19.3 Shaping Based on the Last Mile Wire Rates

Accurate QoS, amongst other things, require that the subscriber rates in the first mile on an MLPPPoX bundle be properly represented in the LNS. In other words, the rate limiting functions in the LNS must account for the last mile on-the-wire encapsulation overhead. The last mile encapsulation can be Ethernet or ATM.

For ATM in the last mile, the LNS accounts for the following per fragment overhead:

- PID
- MLPPP encapsulation header
- ATM Fixed overhead (ATM encap + fixed AAL5 trailer)
- 48B boundary padding as part of AAL5 trailer
- 5B per each 48B of data in ATM cell.

In case of Ethernet encapsulation in the last mile, the overhead is:

- PID
- MLPPP header per fragment
- Ethernet Header + FCS per fragment
- Preamble + IPG overhead per fragment

The **encap-offset** command under the sub-profile>egress CLI node is ignored in case of MLPPPoX. MLPPPoX rate calculation is, by default, always based on the last-mile wire overhead.

The HQoS rates (port-scheduler, aggregate-rate-limit and scheduler) on LNS are based on the wire overhead of the entity to which the HQoS is applied. For example, if the port-scheduler is managing bandwidth on the link between the BB-ISA and the carrier IOM, then the rate of such scheduler accounts for the q-in-q Ethernet encapsulation on that link along with the preamble and inter packet gap (20B).

Figure 43 **QoS Enforcement Points in the LNS**

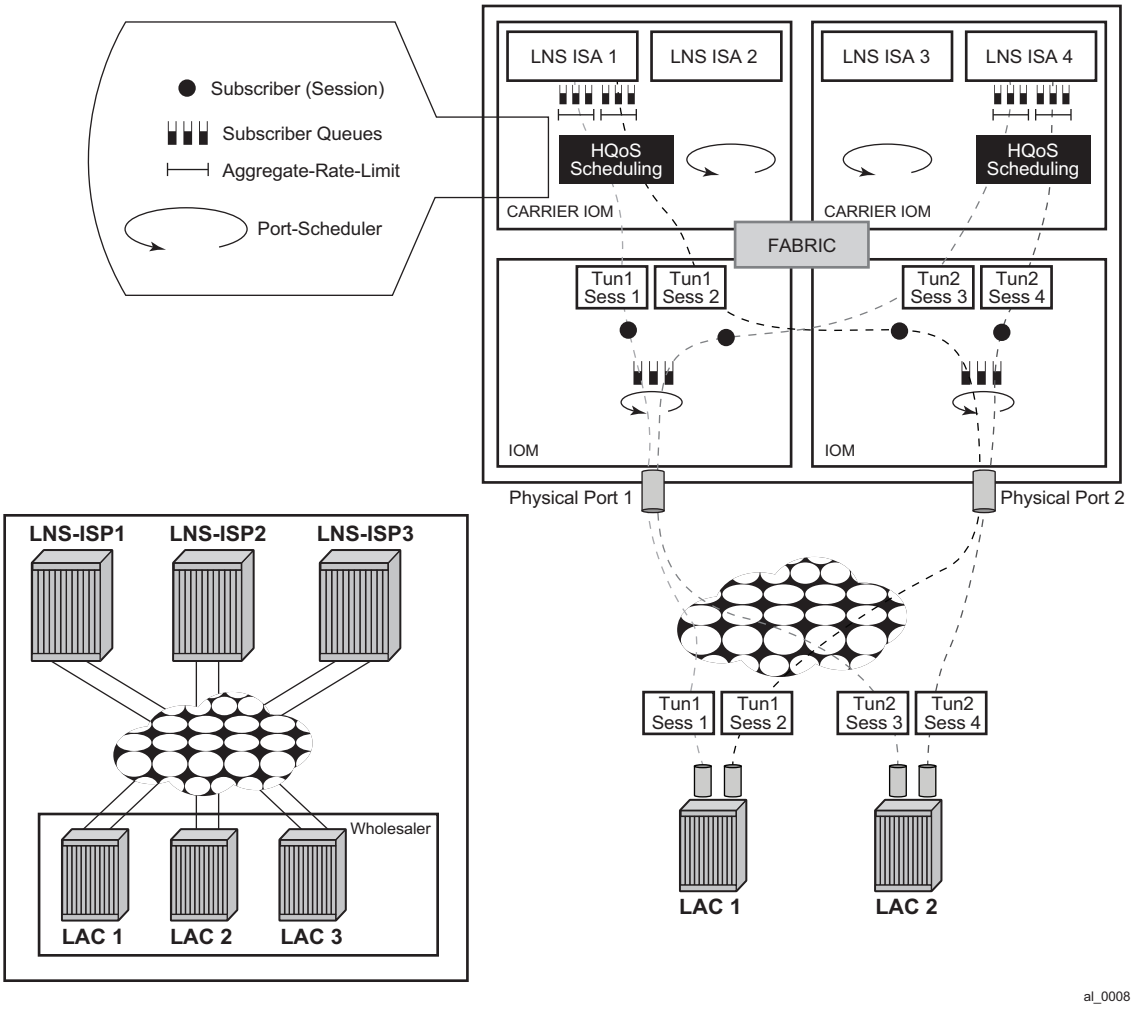

While virtual schedulers (attached via sub-profile) are supported on LNS for plain PPPoX sessions, they are not supported for MLPPPoX bundles. Only aggregate-rate-limit along with the port-scheduler can be used in MLPPPoX deployments.

5.2.19.4 Downstream Bandwidth Management on Egress Port

Bandwidth management on the egress physical ports (Physical Port 1 and Physical Port 2 in Figure 8) is performed at the egress port itself on the egress IOM instead on the carrier IOM. By default, the forwarding class (FC) information is preserved from network ingress to network egress. However, this can be changed via QoS configuration applied to the egress SAP of the carrier IOM towards the BB-ISA.

L2TP traffic originated locally in LNS can be marked via the router/service vprn->sgt-qos hierarchy.

5.2.20 Sub/Sla-Profile Considerations

Sub-profile

In the MLPPPoX case on LNS, multiple sessions are tied into the same subscriber aggregate-rate-limit via a sub-profile. The consequence is that the aggregate rate of the subscriber can be adjusted dynamically depending on the advertised link speed in the last mile and the number of links in the bundle.

Note: Shaping in the LNS is performed per the entire MLPPPoX bundle (subscriber) rather than per individual member links within the bundle. The exception is obviously a MLPPPoX bundle with the single member link (interleaving case) where the relationship between the session and the MLPPPoX bundle is 1:1.

In the LAC, the subscriber aggregate rate cannot be dynamically changed based on the number of links in the bundle and their rate. The LAC has no notion of MLPPPoX bundles. However, multiple sessions that in reality belong to an MLPPPoX bundle under the subscriber are shaped as an aggregate (agg-rate-limit under the sub-profile). This in essence yields the same shaping behavior as on LNS.

Sla-profile

Sessions within the MLPPPoX bundle in LNS share a single sla-profile instances (queues).

In the LAC, as long as the sessions within the subscriber6 are on the same SAP, they can also share the same sla-profile. This is be the case in MLPPPoX.

The manner in which sub/sla-profile are applied to MLPPPoX bundles and the individual sessions within results in aggregate shaping per MLPPPoX bundle as well as allocation of unique set of queues per MLPPPoX bundle. This is valid irrespective of the location where shaping is executed (LAC or LNS). Other vendors may have implemented shaping per session within the bundle and this is something that needs to be taken into consideration during the migration process.

5.2.21 Example of MLPPPoX Session Setup Flow

LAC behavior

- A new PPP(oEoA) session request arrives to the LAC (PADI or LCP Conf Req).
- The LAC negotiates PADx session if applicable.
- The LAC may negotiate MLPPPoX LCP phase with its own endpoint discriminator, or it may reject MLPPPoX specific options in LCP if MLPPPoX on the LAC is disabled (such as, no accept-mrru in the LAC's ppp-policy). If MLPPPoX options (seq num header format, ED, MRRU) are rejected, the assumption is that the client renegotiates plain PPP(oEoA) session with the LAC.
- Once LCP (MLPPPoX capable or not) is negotiated, the session is authenticated (PAP/CHAP).
- Upon successful authentication, an L2TP tunnel is identified to which the session belongs.
- If the session is a non-L2TP session (PTA MLPPPoX capable session for which the tunnel cannot be determined), the session is terminated.
- Otherwise, the QoS constructs are created for the subscriber hosts: the session is assigned to a sub/sla-profiles.
- The session LCP parameters are sent to the LNS via call management messages.

Note: If another LCP session is requested on the same bundle, the LAC creates a new LCP session and join this session to the existing subscriber as another host. In other words, the LAC is bundle agnostic and the two sessions appears as two hosts under the same subscriber.

The following assumes that MLPPPoX is configured on the LNS under the L2TP group or the tunnel hierarchy.

LNS behavior

- The LNS have the option to accept the LCP parameters or to reject them and start renegotiating LCP parameters directly with the client.
- If the LNS choose to renegotiate LCP parameters with the client directly, this renegotiation is completely transparent to the LAC by the means of a T-bit (control vs. data) in the L2TP header. LCP is renegotiated on the LNS with all the options necessary to support MLPPPoX.

Note: Endpoint Discriminator is not mandatory in the MLPPPoX negotiation. If the client rejects it, the LNS must still be able to negotiate MLPPPoX capable session (same is valid for the LAC). If the client's endpoint discriminator is invalid (bad format, invalid class, and so on), the 7750 SR is not negotiate MLPPPoX and instead a plain PPP session is created.

- If the LNS is configured to accept the LCP Proxy parameters, the LNS determines the capability of the client.

If there is no indication of MLPPPoX capability in the Proxy LCP (not even in the original ConfReq), the LNS may accept plain (non MLPPPoX capable) LCP session or renegotiate from scratch the non MLPPPoX capable session.

If there is an indication of MLPPPoX capability in the Proxy LCP (either completely negotiated on the LAC or at least attempted from the client), the LNS tries to accept the MLPPPoX negotiated session by the LAC or renegotiate the MLPPPoX capable session directly with the client.

If the LCP Proxy parameters with MLPPPoX capability are accepted by the LNS then the endpoint as negotiated on the LAC is also accepted.

- Once the MLPPPoX capable LCP session is negotiated or accepted, authentication can be performed on the LNS. Authentication on the LNS can be restarted (CHAP challenge/response with the client), or accepted (chap challenge/response accepted and verified by the LNS via RADIUS).

Note: chap-challenge length is configurable in LNS.

- If the authentication is successful, depending on the evaluation of the parameters negotiated up to this point a new MLPPPoX bundle is created or an existing MLPPPoX bundle is joined. In case that a new bundle is established, the QoS constructs for the subscriber(-host) is created (sub/sla-profile). Session negotiation advances to IPCP phase.
- The decision whether a new session should join an existing MLPPPoX bundle, or trigger creation of a new one is governed by RFC 1990, *The PPP Multilink Protocol (MP)*, section 5.1.3, page 16, cases 1,2,3, and 4.

Note: Interleaving is supported only on MLPPPoX bundles with single session in them.

5.2.22 Other Considerations

- IPv6 is supported.
- AA is supported at LNS where full IP packets can be redirected via AA policies.
- Intra-chassis redundancy is supported:
  - CPM stateful failover
  - BB-ISA — non-stateful failover

5.3 Configuration Notes

MLPPP in subscriber management context is supported only over ATM, Ethernet over ATM or plain Ethernet transport (MLPPPoX). Native MLPPP over PPP/HDLC links is supported outside of the subscriber management context on the ASAP MDA.

MLPPPoX is supported only on LNS.

Interleaving is supported only on MLPPPoX bundles with a single member link. If more than one link is present in an MLPPPoX bundle, the interleaving is automatically disabled and a SNMP trap is generated. The MIB for this even is defined as tmnxMlpppBundleIndicatorsChange.

If MLPPPoX is enabled on LNS, the load balancing mode between the BB-ISAs within the group should be set to per tunnel. This ensures that all sessions of the same MLPPPoX bundle are terminated on the same BB-ISA. On the LAC, sessions of the same bundle are setup in the same tunnel.

Virtual schedulers are not supported on MLPPPoX tunnels on LNS. However, aggregate-rate-limit is supported.

The aggregate-rate-limit on LNS is automatically adjusted to the minimum value of:

- configured aggregate-rate-limit
- minimum last mile rate (obtained via LUDB, RADIUS or PPPoE tags) multiplied by the number of links in the bundle.

The aggregate-rate-limit on the LAC is not adjusted automatically. Therefore, if configured it should be set to a high value and thus the traffic treatment should rely on QoS performed on the LNS.

The rate (rate-down information) of the member links within the bundle must be the same. Otherwise the lowest rate is selected and applied to all member links.

A single CoA for a rate change (Alc-Access-Loop-Rate-Down) of an individual link in an MLPPPoX bundle modifies rates of all links in the bundle. This is applicable on LNS only.

The range of supported last mile rate (rate-down information) for the member links on an MLPPPoX session is 1 kb/s — 100 Mb/s. On the LNS the last mile rate can be obtained:

- From the LAC via Tx-Connect-Speed AVP or by standard L2TP encoding as described in the RFC 5515, *Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extensions.*

- From the LAC via LUDB or RADIUS
- Directly on the LNS via LUDB or RADIUS.

The session fails to come up if the obtained rate-down information is outside of the allowable range (1 kb/s — 100 Mb/s).

A session within the MLPPPoX bundle is terminated if the rate-down information for the session is out of bounds (1 kb/s — 100 Mb/s).

If a member link in the last mile fails, traffic is blackholed until the LNS is notified of this failure. The failure detection in the LNS relies on PPP keepalives.

Shaping is performed per MLPPPoX bundle and not individually per member links.

If encapsulation overhead associated with fragmentation is too large in comparison to payload, the fragments are sized based on the encapsulation overhead (to increase link efficiency) rather than on maximum transmission delay.

There can be only a single MLPPPoX bundle per subscriber.

MLPPPoX bundles and non-MLPPPoX (plain L2TP PPPoE) sessions cannot coexist under the same subscriber.

Filters and mirrors (LI) are not supported on MLPPPoX bundles on LAC.

ip-only type mirrors are supported on MLPPPoX bundles.

In MLPPP scenario, downstream traffic is traversing Carrier IOM and BB-ISA twice. This is referred to as dual-pass and effectively cuts the throughput for MLPPP in half (for example, 5Gb/s of MLPPP traffic on a 10Gb/s capable BB-ISA).

5.4 PPP Command Reference

5.4.1 Command Hierarchies

- PPPoE Policy Configuration Commands
- PPPoE Service Commands
- PPPoE Local User Database Commands
- MLPPP on LNS Commands
- Show Commands

5.4.1.1 PPPoE Policy Configuration Commands

```
config
    — subscriber-mgmt
        — ppp-policy ppp-policy-name [create]
        — no ppp-policy ppp-policy-name
            — default-pap-password password [hash | hash2]
            — no default-pap-password
            — default-user-name ppp-username
            — no default-user-name
            — description description-string
            — no description
            — [no] disable-cookies
            — [no] force-ppp-mtu-gt-1492
            — [no] ipcp-subnet-negotiation
            — keepalive seconds [hold-up-multiplier multiplier]
            — no keepalive
            — [no] lcp-ignore-magic-numbers
            — max-sessions-per-mac sessions [allow-same-circuit-id-for-dhcp]
            — no max-sessions-per-mac
            — pado-ac-name name
            — no pado-ac-name
            — pado-delay deci-seconds
            — no pado-delay
            — ppp-authentication {pap | chap | pref-chap | pref-pap}
            — no ppp-authentication
            — ppp-chap-challenge-length min minimum-length max maximum-length
            — no ppp-chap-challenge-length
            — [no] ppp-initial-delay
            — ppp-mtu mtu-bytes
            — no ppp-mtu
            — ppp-options
                — custom-option protocol option-number address ip-address
                — custom-option protocol option-number hex hex-string
```

```
            — custom-option protocol option-number  string ascii-string
            — no custom-option protocol option-number
        — re-establish-session padr
        — no re-establish-session
        — [no] reject-disabled-ncp
        — [no] reply-on-padt
        — session-timeout timeout
        — no session-timeout
        — unique-sid-per-sap [per-msap]
        — no unique-sid-per-sap
```

5.4.1.2 PPPoE Service Commands

Note: The commands listed in this section apply only to the 7750 SR.

```
config
    — service
        — ies service-id [customer customer-id] [vpn vpn-id] [create]
        — vprn service-id [customer customer-id] [create]
        — no vprn service-id
        — no ies service-id
            — [no] subscriber-interface ip-int-name
                — [no] group-interface ip-int-name
                    — dhcp
                        — client-applications {[dhcp] [ppp]}
                        — no client-applications
                    — [no] ppp
                        — description description-string
                        — no description
                        — policy ppp-policy-name
                        — no policy
                        — session-limit session-limit
                        — no session-limit
                        — [no] shutdown
                        — user-db local-user-db-name
                        — no user-db
                    — [no] pppoe
                        — anti-spoof pppoe-anti-spoofing-type
                        — no anti-spoof
                        — description description-string
                        — no description
                        — dhcp-client
                            — [no] ccag-use-origin-sap
                        — policy ppp-policy-name
                        — no policy
                        — python-policy name
                        — no python-policy
```

```
— sap-session-limit sap-session-limit
— no sap-session-limit
— session-limit session-limit
— no session-limit
— [no] shutdown
— user-db local-user-db-name
— no user-db
```

5.4.1.3 PPPoE Local User Database Commands

Note: The commands listed in this section apply only to the 7750 SR only.

```
config
    — subscriber-mgmt
        — local-user-db local-user-db-name [create]
        — no local-user-db local-user-db-name
            — ppp
                — mask {[prefix-string prefix-string | prefix-length prefix-length] [suffix-
                    string suffix-string | suffix-length suffix-length]}
                — no mask
                — host host-name [create]
                — no host host-name
                    — acct-policy acct-policy-name [duplicate acct-policy-name]
                    — no acct-policy
                    — address gi-address
                    — address ip-address
                    — address pool pool-name
                    — no address
                    — authentication-policy policy-name
                    — no authentication-policy
                    — host-identification
                        — circuit-id string ascii-string
                        — circuit-id hex hex-string
                        — no circuit-id
                        — mac ieee-address
                        — no mac
                        — remote-id remote-id
                        — no remote-id
                        — service-name service-name
                        — no service-name
                        — username user-name [no-domain]
                        — username user-name domain-only
                        — no username
                    — identification-strings option-number [create]
                    — no identification-strings
                        — ancp-string ancp-string
                        — no ancp-string
```

```
                        — app-profile-string app-profile-string
                        — no app-profile-string
                        — inter-dest-id intermediate-destination-id
                        — no inter-dest-id
                        — sla-profile-string sla-profile-string
                        — no sla-profile-string
                        — spi-sharing-group-id group-id
                        — no spi-sharing-group-id
                        — sub-profile-string sub-profile-string
                        — no sub-profile-string
                        — subscriber-id sub-ident-string
                        — no subscriber-id
                    — l2tp
                        — group tunnel-group-name
                        — no group
                    — options
                        — custom-option option-number address [ip-address]
                        — custom-option option-number hex hex-string
                        — custom-option option-number string ascii-string
                        — no custom-option option-number
                        — dns-server [ip-address]
                        — no dns-server
                        — netbios-name-server ip-address [ip-address]
                        — no netbios-name-server
                    — pado-delay deci-seconds
                    — no pado-delay
                    — password {ignore | chap string | pap string}
                    — no password
                    — retail-service-id service-id
                    — no retail-service-id
                    — [no] shutdown
                — mask type pppoe-match-type {[prefix-string prefix-string | prefix-
                    length prefix-length] [suffix-string suffix-string | suffix-length
                    suffix-length]}
                — no mask type pppoe-match-type
                — match-list pppoe-match-type-1 [pppoe-match-type-2]
                — no match-list
            — [no] shutdown
```

5.4.1.4 MLPPP on LNS Commands

Refer to the *7450 ESS, 7750 SR, and VSR Multiservice Integrated Service Adapter Guide* for MLPPP configuration and command information.

```
config
    — subscriber-mgmt
        — ppp-policy ppp-policy-name [create]
        — no ppp-policy ppp-policy-name
            — mlppp
                — [no] accept-mrru
```

```
                        — endpoint ip ip-address
                        — endpoint mac ieee-address
                        — endpoint system-ip
                        — endpoint system-mac
                        — no endpoint
                        — [no] short-sequence-numbers
                    — local-user-db local-user-db-name [create]
                    — no local-user-db local-user-db-name
                        — ppp
                            — host host-name [create]
                            — no host host-name
                                — [no] access-loop
                                    — encap-offset [type encap-type]
                                    — no encap-offset
                                    — rate-down rate
                                    — no rate-down

config
    — router
        — l2tp
            — group tunnel-group-name [create]
            — no group tunnel-group-name
                — load-balance-method {session | tunnel}
                — no load-balance-method
                — mlppp
                    — endpoint ip ip-address
                    — endpoint mac ieee-address
                    — endpoint system-ip
                    — endpoint system-mac
                    — no endpoint
                    — [no] interleave
                    — max-fragment-delay milliseconds
                    — no max-fragment-delay
                    — max-link max-links
                    — no max-link
                    — reassembly-timeout {{100 | 1000} milliseconds}
                    — no reassembly-timeout
                — tunnel tunnel-name [create]
                — no tunnel tunnel-name
                    — load-balance-method {session | tunnel}
                    — no load-balance-method
                    — mlppp
                        — admin-state {up | down}
                        — no admin-state
                        — endpoint ip ip-address
                        — endpoint mac ieee-address
                        — endpoint system-ip
                        — endpoint system-mac
                        — no endpoint
                        — interleave {always | never}
                        — no interleave
                        — max-fragment-delay milliseconds
                        — no max-fragment-delay
                        — max-link max-links
```

```
                                        — no max-link
                                        — reassembly-timeout {{100 | 1000} milliseconds}
                                        — no reassembly-timeout

config
    — service
        — vprn
            — l2tp
                — group
                    — load-balance-method {session | tunnel}
                    — no load-balance-method
                    — mlppp
                        — admin-state {up | down}
                        — no admin-state
                        — endpoint ip ip-address
                        — endpoint mac ieee-address
                        — endpoint system-ip
                        — endpoint system-mac
                        — no endpoint
                        — interleave {always | never}
                        — no interleave
                        — max-fragment-delay milliseconds
                        — no max-fragment-delay
                        — max-link max-links
                        — no max-link
                        — reassembly-timeout {{100 | 1000} milliseconds}
                        — no reassembly-timeout
                    — tunnel
                            — load-balance-method {session | tunnel}
                            — no load-balance-method
                        — mlppp
                            — admin-state {up | down}
                            — no admin-state
                            — endpoint ip ip-address
                            — endpoint mac ieee-address
                            — endpoint system-ip
                            — endpoint system-mac
                            — no endpoint
                            — interleave {always | never}
                            — no interleave
                            — load-balance-method {session | tunnel}
                            — no load-balance-method
                            — max-fragment-delay milliseconds
                            — no max-fragment-delay
                            — max-link max-links
                            — no max-link
                            — reassembly-timeout {{100 | 1000} milliseconds}
                            — no reassembly-timeout
```

5.4.1.5 Show Commands

```
show
    — router
        — l2tp
            — peer ip-address [udp-port port]
            — peer ip-address statistics [udp-port port]
            — peer [draining] [blacklisted | selectable | unreachable]
    — service
        — id service-id
            — ipoe
                — session [sap sap-id] [mac ieee-address] [circuit-id circuit-id] [remote-id remote-id] [interface ip-int-name | ip-address] [inter-dest-id intermediate-destination-id] [no-inter-dest-id] [ip-address ip-prefix[/prefix-length]] [port port-id] [subscriber sub-ident-string] [sap-session-id sap-session-index] [detail] [wholesaler service-id] [router-advertisement-policy policy-name]
                — summary
            — pppoe
                — session [interface ip-int-name | ip-address | sap sap-id] [type pppoe-session-type] [session-id session-id] [mac ieee-address] [ip-address ip-address[/mask]] [port port-id] [no-inter-dest-id | inter-dest-id intermediate-destination-id] [steering-profile steering-profile] [router-advertisement-policy policy-name] [detail | statistics]
                — session l2tp-connection-id connection-id [detail | statistics]
                — statistics [{sap sap-id | interface ip-int-name | ip-address}
                — summary
    — subscriber-mgmt
        — ppp-policy [ppp-policy-name [association]]
```

5.4.2 Command Descriptions

5.4.2.1 Generic Commands

description

Syntax **description** *description-string*
no description

Context config>subscr-mgmt>pppoe-policy
config>service>ies>sub-if>grp-if>pppoe
config>service>vprn>sub-if>grp-if>pppoe

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Parameters *description-string* — Specifies the description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

shutdown

Syntax [**no**] **shutdown**

Context config>service>ies>sub-if>grp-if>pppoe
config>service>vprn>sub-if>grp-if>pppoe

Description This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command places the entity into an administratively enabled state.

5.4.2.1.1 Subscriber Management Commands

subscriber-mgmt

Syntax **subscriber-mgmt**

Context config>subscr-mgmt

Description This command enables the context to configure subscriber management parameters.

local-user-db

Syntax **local-user-db** *local-user-db-name* [**create**]
no local-user-db *local-user-db-name*

Context config>subscr-mgmt

Description This command enables the context to configure a local user database.

The **no** form of this command reverts to the default.

Parameters *local-user-db-name* — Specifies the name of a local user database.

ppp

Syntax **ppp**

Context config>subscr-mgmt>loc-user-db

Description This command enables the context to configure PPP host parameters.

mask

Syntax **mask type** *dhcp-match-type* {[**prefix-string** *prefix-string* | **prefix-length** *prefix-length*] [**suffix-string** *suffix-string* | **suffix-length** *suffix-length*]}
no mask type *dhcp-match-type*

Context config>subscr-mgmt>loc-user-db>ipoe
config>subscr-mgmt>loc-user-db>ppp
config>subscr-mgmt>loc-user-db>ipoe

Description This command configures the mask.

The **no** form of this command removes the parameters from the mask configuration.

Parameters *dhcp-match-type* — Specifies up to four matching types to identify a host.

Values DHCP: circuit-id, option60, remote-id, sap-id, string, system-id
PPP: circuit-id, remote-id, service-name, username

Values prefix-string *prefix-string*

Specifies a substring that is stripped of the start of the incoming circuit ID before it is matched against the value configured in the DHCP or PPPoE circuit ID.

This string can only contain printable ASCII characters. The “*” character is a wildcard that matches any substring. If a “\” character is masked, use the escape key so it becomes “\\”.

Values 127 characters maximum, *' is a wildcard

prefix-length *prefix-length* — Specifies the number of characters to remove from the start of the incoming circuit ID before it is matched against the value configured in the DHCP circuit ID.

Values 1 to 127

suffix-string *suffix-string* — Specifies a substring that is stripped of the end of the incoming circuit ID before it is matched against the value configured in DHCP circuit ID.

This string can only contain printable ASCII characters. The "*" character is a wildcard that matches any substring. If a "\" character is masked, use the escape key so it becomes "\\".

Values 127 characters maximum

suffix-length *suffix-length* — Specifies the number of characters to remove from the end of the incoming circuit ID before it is matched against the value configured in the DHCP circuit ID.

Values 1 to 127

host

Syntax **host** *host-name* [**create**]
no host *host-name*

Context config>subscr-mgmt>loc-user-db>ipoe
config>subscr-mgmt>loc-user-db>ppp

Description This command defines a DHCP or PPP subscriber.

The **no** form of this command removes the host name from the configuration.

Parameters *host-name* — Specifies the host name.

create — Keyword used to create the host name. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

ppp-policy

Syntax **ppp-policy** *ppp-policy-name* [**create**]
no ppp-policy *ppp-policy-name*

Context config>subscr-mgmt

Description This command configures a PPP policy. These policies are referenced from interfaces configured for PPP. Multiple PPP policies may be configured.

This default policy cannot be modified nor deleted.

Default default

Parameters *ppp-policy-name* — Specifies the PPP policy name up to 32 characters.

create — Keyword used to create the entity. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

default-pap-password

Syntax **default-pap-password** *password* [**hash** | **hash2**]
no default-pap-password

Context config>subscr-mgmt>ppp-policy

Description This command configures the default PAP password for RADIUS authentication when the Password-Length=0 in the PAP Authenticate-Request.

RADIUS authentication cannot be initiated when the Password-Length=0 in the PAP Authenticate-Request and no default-pap-password is configured. The PPP session terminates in this case.

Parameters *password* — Specifies a default PAP password up to 64 characters.

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

default-user-name

Syntax **default-user-name** *ppp-username*
no default-user-name

Context config>subscr-mgmt>ppp-policy

Description This command configures the default username for authentication when not provided in PAP/CHAP authentication (no Name field in CHAP Response message or Peer-Id-Length=0 in PAP Authenticate-Request).

The PPP session terminates when no username is provided in PAP/CHAP authentication and no default-user-name is configured.

Parameters *ppp-username* — Specifies a default username up to 253 characters.

disable-cookies

Syntax [**no**] **disable-cookies**

Context config>subscr-mgmt>ppp-policy

Description This command disables the use of cookies.

The **no** form of this command enables cookies.

force-ppp-mtu-gt-1492

Syntax [**no**] **force-ppp-mtu-gt-1492**

Context config>subscr-mgmt>ppp-policy

Description This command enables PPPoE Maximum-Receive-Unit (MRU) negotiations greater than 1492 bytes without the need to receive a “PPP-Max-Payload” tag in PADI/PADR from the client as defined in RFC 4638, *Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE)*.

The MRU send in the initial LCP Config Request is determined by the **port mtu** and **ppp-policy ppp-mtu** parameters.

keepalive

Syntax **keepalive** *seconds* [**hold-up-multiplier** *multiplier*]
no keepalive

Context config>subscr-mgmt>ppp-policy

Description This command defines the keepalive interval and the number of keepalives that can be missed before the session is declared down for this PPP policy.

The **no** form of this command reverts to the default value.

Default keepalive 30 hold-up-multiplier 3

Parameters *seconds* — Specifies the keepalive interval in seconds.

Values 10 to 300

hold-up-multiplier *multiplier* — Specifies the number of keepalives that can be missed.

Values 1 to 5

ipcp-subnet-negotiation

Syntax [**no**] **ipcp-subnet-negotiation**

Context config>subscr-mgmt>ppp-policy

Description This command enables subnet negotiation using PPP IPCP Subnet-Mask option (0x90) if requested by the client. The subnet can be obtained from RADIUS (Framed-IP-Netmask attribute) or local user database. The subnet is installed as a managed route of the PPP session. This requires the anti-spoof type on the SAP to be configured to nh-mac.

By default, an IPCP Config Request with IPCP Subnet-Mask option (0x90) is rejected.

lcp-ignore-magic-numbers

Syntax [**no**] **lcp-ignore-magic-numbers**

Context config>subscr-mgmt>ppp-policy

Description This command enables the PPP session to stay established when an LCP peer magic number mismatch is detected.

By default, the PPP session is terminated when an LCP peer magic number mismatch is detected.

max-sessions-per-mac

Syntax **max-sessions-per-mac** *sessions* [**allow-same-circuit-id-for-dhcp**]
no max-sessions-per-mac

Context config>subscr-mgmt>ppp-policy

Description This command sets the maximum PPP sessions that can be opened for a given MAC address.

To enable IPv4 address allocation using the internal DCHCPv4 client for multiple PPPoE sessions on a single SAP and having the same MAC address and circuit-ID, the optional CLI parameter **allow-same-circuit-id-for-dhcp** should be added. The SR OS local DHCP server detects the additional vendor-specific options inserted by the internal DCHCPv4 client and use an extended unique key for lease allocation.

The **no** form of this command reverts to the default value.

Parameters *sessions* — Specifies the maximum PPP sessions that can be opened for the given MAC address.

Values 1 to 8191

allow-same-circuit-id-for-dhcp — Enables support for IPv4 address allocation using the internal DCHCPv4 client for multiple PPPoE sessions on a single SAP that have the same MAC address and circuit ID.

pado-ac-name

Syntax **pado-ac-name** *name*
no pado-ac-name

Context config>subscr-mgmt>ppp-policy

Description This command configures the Access Concentrator name that is used in the PPPoE PADO message.

By default, the system name or if not configured, the chassis Serial Number is used.

Parameters *name* — Specifies the string up to 128 characters to be used as AC name in the PPPoE PADO message.

pado-delay

Syntax **pado-delay** *deci-seconds*
no pado-delay

Context config>subscr-mgmt>ppp-policy

Description This command configures the delay timeout before sending a PPP Active Discovery Offer (PADO) packet.

Parameters *deci-seconds* — Specifies the delay timeout before sending a PADO.

Values 1 to 30

ppp-authentication

Syntax **ppp-authentication** {**pap** | **chap** | **pref-chap** | **pref-pap**}
no ppp-authentication

Context config>subscr-mgmt>ppp-policy

Description This command configures the PPP protocol used to authenticate the PPP session.

Default ppp-authentication pref-chap

Parameters **pap** — Specifies to always use PAP to authenticate the sessions.

chap — Specifies to always use CHAP to authenticate the sessions.

pref-chap — Specifies to attempt to use CHAP and if it fails, use PAP.

pref-pap — Specifies to attempt to use PAP and if it fails, use CHAP.

ppp-chap-challenge-length

Syntax **ppp-chap-challenge-length min** *minimum-length* **max** *maximum-length*
no ppp-chap-challenge-length

Context config>subscr-mgmt>ppp-policy

Description This command configures the minimum and maximum length of a PPP Chap Challenge.

When the Chap Challenge is exactly 16 bytes, it is send in the [60] CHAP-Challenge RADIUS attribute and also copied in the RADIUS Authenticator field from the RADIUS Access Request.

Default ppp-chap-challenge-length min 32 max 64

Parameters **min** *minimum-length* — Specifies the minimum PPP CHAP challenge length.

Values 8 to 64

max *maximum-length* — Specifies the maximum PPP CHAP challenge length.

Values 8 to 64

ppp-initial-delay

Syntax [**no**] **ppp-initial-delay**

Context config>subscr-mgmt>ppp-policy

Description This command delays the sending of an LCP-configure request after the discovery phase by 40 – 60 milliseconds.

ppp-mtu

Syntax **ppp-mtu** *mtu-bytes*
no ppp-mtu

Context config>subscr-mgmt>ppp-policy

Description This command configures the maximum PPP MTU size.

Parameters *mtu-bytes* — Specifies the maximum PPP MTU size.

Values 512 to 9212

ppp-options

Syntax **ppp-options**

Context config>subscr-mgmt>ppp-policy

Description This command enables the context to configure PPP options.

custom-option

Syntax **custom-option** *protocol option-number* **address** *ip-address*
custom-option p*rotocol option-number* **hex** *hex-string*
custom-option *protocol option-number* **string** *ascii-string*
no custom-option *protocol option-number*

Context config>subscr-mgmt>ppp-policy>options

Description This command provides the ability to configure custom PPP options.

Note: Standard options such as the DNS name is returned from DHCP or RADIUS and be converted to PPP automatically. Compression is not supported.

The **no** form of this command removes the custom options from the configuration.

Parameters *protocol* — Specifies a protocol for the custom option.

Values lcp, ipcp, ipv6cp

option-number — Assigns an identifying number for the custom option.

Values 0 to 255

ip-address — Specifies the IP address in the a.b.c.d format.

ascii-string — Specifies an ASCII format string for the custom option up to 127 characters.

hex-string — Specifies a hex value for the custom option.

Values [0x0 to 0xFFFFFF (up to 254 hex nibbles)

re-establish-session

Syntax **re-establish-session padr**
no re-establish-session

Context config>subscr-mgmt>ppp-policy

Description This command enables host to reconnect and override existing session.

If disabled and a subscriber abruptly terminates a PPP sessions without sending a PADT to the BNG, the BNG denies any reconnect attempts until the stale PPP session has expired. With this, enabled re-establish-session eliminates the waiting period by allowing immediate PPP reconnection attempts.

The **no** form of this command reverts to the default.

reject-disabled-ncp

Syntax [**no**] **reject-disabled-ncp**

Context config>subscr-mgmt>ppp-policy

Description This command forces an LCP Protocol Reject when receiving an IPv6CP Configure Request message while IPv6 is not configured.

By default, an IPv6CP Configure Request message is silently ignored when IPv6 is not configured.

reply-on-padt

Syntax [**no**] **reply-on-padt**

Context config>subscr-mgmt>ppp-policy

Description This command enables replying to PPPoE Active Discovery Terminate (PADT) packets. Some of the PPPoE clients expect reply on PADT message before the context of the session is cleared up. To support such client, a command enabling reply to PADT is provided.

The **no** form of this command reverts to the default.

session-timeout

Syntax **session-timeout** *timeout*
no session-timeout

Context config>subscr-mgmt>ppp-policy

Description This command defines the time before the PPP session is terminated. T

A RADIUS specified session-timeout (attribute [27] Session-Timeout) overrides the CLI configured value.

The **no** form of this command reverts to the default.

Parameters *timeout* — Specifies the session timeout in seconds.

Values 1 to 31104000

unique-sid-per-sap

Syntax **unique-sid-per-sap** [**per-msap**]
no unique-sid-per-sap

Context config>subscr-mgmt>ppp-policy

Description This command assigns a unique session ID to each PPPoE session with different MAC addresses that are active on a single SAP.

On a capture SAP, a unique session ID is assigned per capture SAP: multiple sessions with different MAC addresses that are active on the same or different MSAP have a unique session ID per capture SAP.

With the optional parameter per-msap, a unique session ID is assigned per MSAP:

- multiple sessions with different MAC addresses that are active on the same MSAP have a unique session ID per MSAP
- multiple sessions with different MAC addresses that are active on different MSAPs are not guaranteed to have a unique session ID

The maximum session ID range is 1 to 8191.

By default, all PPPoE sessions with different MAC address on a given SAP or MSAP have session-id 1.

The **no** form of this command reverts to the default.

Parameters **per-msap** — Assigns a unique session id for each session with different MAC address that is active on the same MSAP. This parameter has no effect on regular SAPs.

5.4.2.1.2 MLPPP on LNS Commands

accept-mrru

Syntax [**no**] **accept-mrru**

Context config>subscr-mgmt>ppp-policy>mlppp

Description This command is applicable only to LAC. MRRU option is an indication that the session is of MLPPPoX type. The 7750 SR LAC never initiates the MRRU option in LCP negotiation process. However, it responds to MRRU negotiation request by the client.

This command provides an option to specifically enable or disable negotiation of MLPPPoX on a capture SAP level or on a group interface level.

The **no** form of this command causes the MRRU option in LCP to not be negotiated by LAC.

admin-state

Syntax **admin-state** {**up** | **down**}
no admin-state

Context config>router>l2tp>group>tunnel>mlppp
config>service>vprn>l2tp>group>tunnel>mlppp

Description This command enables MLPPP for this tunnel group and is applicable only to LNS.

The tunnel can be explicitly activated (assuming that the parent group is in a no shutdown state) or deactivated by the **up** and **down** keywords.

If there the admin state is not configured, the tunnel inherits its administrative state from its parent (group).

The **no** form of this command causes the tunnel administrative state to be inherited from the group.

Parameters **up** — Specifies that the tunnel is to be administratively up.

down — Specifies that the tunnel is to be administratively down.

encap-offset

Syntax **encap-offset** [**type** *encap-type*]
no encap-offset

Context config>subscr-mgmt>loc-user-db>ppp>host>ale

Description This command is applicable within the LAC/LNS context. It provides the last mile link encapsulation information that is needed for proper (shaping) rate calculations and interleaving delay in the last mile.

The encapsulation value is taken from the following sources in the order of priority:

- Statically provisioned value in local user database (LUDB).
- RADIUS
- PPPoE tags on LAC or ICRQ message (RFC 5515) on LNS

If the encapsulation information is not provided by any of the existing means (LUDB, RADIUS, AVP signaling, PPPoE Tags), then by default pppoe-null encapsulation is in effect.

The following values are supported encapsulation values on LNS in the 7750 SR.

encap-type:

pppoa-llc LLC (NLPID) PPPoA encapsulation
pppoa-null VC-MUX PPPoA encapsulation

pppoeoa-llc	LLC/SNAP based bridged Ethernet PPPoEoA encapsulation without FCS
pppoeoa-llc-fcs	LLC/SNAP based bridged Ethernet PPPoEoA encapsulation with FCS
pppoeoa-null	VC-MUX PPPoEoA encapsulation without FCS
pppoeoa-null-fcs	VC-MUX PPPoEoA encapsulation with FCS
pppoe	PPPoE encapsulation
pppoe-tagged	Tagged PPPoE encapsulation

The following are not supported encapsulation values on LNS in the 7750 SR.

- pppoeoa-llc-tagged
- pppoeoa-llc-tagged-fcs
- pppoeoa-null-tagged
- pppoeoa-null-tagged-fcs
- ipoa-llc
- ipoa-null
- ipoeeoa-llc
- ipoeoa-llc-fcs
- ipoeoa-llc-tagged
- ipoeoa-llc-tagged-fcs
- ipoeoa-null
- ipoeoa-null-fcs
- ipoeoa-null-tagged
- ipoeoa-null-tagged-fcs
- ipoe
- ipoe-tagged

The **no** form of this command reverts to the default.

endpoint

Syntax **endpoint ip** *ip-address*
endpoint mac *ieee-address*
endpoint system-ip
endpoint system-mac
no endpoint

Context config>router>l2tp>group>mlppp
config>router>l2tp>group>tunnel>mlppp
config>service>vprn>l2tp>group>mlppp
config>service>vprn>l2tp>group>tunnel>mlppp

config>subscr-mgmt>ppp-plcy>mlppp

Description When configured under the l2tp hierarchy, this command is applicable to LNS.

Within the ppp-policy, this command is applicable only to LAC.

The endpoint, according to RFC 1990, represents the system transmitting the packet. It is used during MLPPPoX negotiation phase to distinguish this peer from all others.

In the case that the client rejects the endpoint option during LCP negotiation, the LAC and the LNS must be able to negotiate the LCP session without the endpoint option.

The **no** form of this command disables sending endpoint option in LCP negotiation.

Parameters *ip-address* — Specifies the IPv4 address (class 2).

system-ip — Specifies to use the system IPv4 address (class 2).

ieee-address — Specifies the MAC address of the interface (class 3).

system-mac — Specifies to use the MAC address of the system (class 3).

interleave

Syntax [**no**] **interleave**

Context config>router>l2tp>group>mlppp
config>service>vprn>l2tp>group>mlppp

Description This command is applicable only to LNS. Interleaving is supported only on MLPPPoX bundles that contain a single member link. If more than one link is present in the MLPPPoX bundle, interleaving is automatically disabled and a TRAP/log (tmnxMlpppBundleIndicatorsChange) is generated.

The minimum supported rate of the link on which interleaving is performed is 1 kb/s.

If configured at this level, interleaving is enabled on all tunnels within the group, unless it is explicitly disable per tunnel.

The **no** form of this command reverts to the default.

interleave

Syntax **interleave** {**always** | **never**}
no interleave

Context config>router>l2tp>group>tunnel>mlppp
config>service>vprn>l2tp>group>tunnel>mlppp

Description This command configures the user of link fragmentation and interleaving and is applicable only to LNS. Interleaving is supported only on MLPPPoX bundles that contain a single member link. If more than one link is present in the MLPPPoX bundle, interleaving is automatically disabled and a TRAP/log (tmnxMlpppBundleIndicatorsChange) is generated.

The minimum supported rate of the link on which interleaving is performed is 1 kb/s.

Interleaving configured on this level overwrites the configuration option under the group hierarchy. If the **no** form of this command is configured for interleaving at this level, the interleaving configuration inherits the configuration option configured under the L2TP group.

The **no** form of this command reverts to the default.

Parameters **always** — Always perform interleaving on single linked MLPPPoX sessions within this tunnel, regardless of the configuration option for interleaving under the group level.

never — Never perform interleaving on single linked MLPPPoX sessions within this tunnel, regardless of the configuration option for interleaving under the group level.

load-balance-method

Syntax **load-balance-method** {**per-session** | **per-tunnel**}
no load-balance-method

Context config>router>l2tp>group
config>router>l2tp>group>tunnel
config>service>vprn>l2tp>group
config>service>vprn>l2tp>group>tunnel

Description This command is applicable only to LNS. By default traffic load balancing between the BB-ISAs is based on sessions. Each session is individually assigned to an BB-ISA during session establishment phase.

By introducing MLPPPoX, all sessions of a bundle must be terminated on the same LNS BB-ISA. This is necessary for two reasons:

- QoS in the carrier IOM has a uniform view of the subscriber
- a single BB-ISA is responsible for MLPPPoX encapsulation/fragmentation for a given bundle.

Therefore, if fragmentation is enabled, load-balancing per tunnel must be configured. In the per tunnel load-balancing mode, all sessions within the same tunnel are terminated on the same LNS BB-ISA.

In the case that we have MLPPPoX sessions with a single member link, both load-balancing methods are valid.

The **no** form of this command reverts to the default.

Default load-balance-method per-session

Parameters **per-session** — The traffic load balancing between the LNS BB-ISAs is based on individual PPPoE sessions.

per-tunnel — The traffic load balancing between the LNS BB-ISAs is based on tunnels.

max-fragment-delay

Syntax **max-fragment-delay** *delay*
no max-fragment-delay

Context config>router>l2tp>group>mlppp
config>router>l2tp>group>tunnel>mlppp
config>service>vprn>l2tp>group>mlppp
config>service>vprn>l2tp>group>tunnel>mlppp

Description This command is applicable only to LNS. It determines the maximum fragment delay caused by the transmission that is imposed on a link.

Fragmentation can be used to interleave high priority packet in-between low priority fragments on a MLPPPoX session with a single link or on a MLPPPoX session with multiple links to better load balance traffic over multiple member links.

Parameters *delay* — Specifies the interval in milliseconds.

Values 5 to 1000ms

max-link

Syntax **max-links** *max-links*
no max-links

Context config>router>l2tp>group>mlppp
config>router>l2tp>group>tunnel>mlppp
config>service>vprn>l2tp>group>mlppp
config>service>vprn>l2tp>group>tunnel>mlppp

Description This command is applicable only to LNS. It determines the maximum number of links that can be put in a bundle.

Any attempt of a session to join a bundle that is above the max-link limit is rejected.

If interleaving is configured, it is recommended that max-links be set to 1 or a version of the command is used (no max-links). Both have the same effect.

The configuration under the tunnel hierarchy overrides the configuration under the group hierarchy.

The **no** form of this command reverts to the default.

Default max-links 1

Parameters *max-links* — Specifies the maximum number of links in a bundle.

Values 1 to 8

reassembly-timeout

Syntax **reassembly-timeout** {*timeout*}
no reassembly-timeout

Context config>router>l2tp>group>mlppp
config>router>l2tp>group>tunnel>mlppp
config>service>vprn>l2tp>group>mlppp
config>service>vprn>l2tp>group>tunnel>mlppp

Description This command is applicable only to LNS. It determines the time during which the LNS keeps fragments of the same packet in the buffer before it discards them. The assumption is that if the fragments do not arrive within certain time, the chance is that they were lost somewhere in the network. In this case the partial packet cannot be reassembled and all fragments that has arrived up to this point and are stored in the buffer IS discarded in order to free up the buffer. Otherwise, a condition arises in which partial packets are held in the buffer until the buffer is exhausted.

The configuration under the tunnel hierarchy overrides the configuration under the group hierarchy.

The **no** form of this command reverts to the default.

Default reassembly-timeout 1000

Parameters *timeout* — Specifies the reassembly timeout value.

Values 100, 1000 milliseconds

rate-down

Syntax **rate-down** *rate*
no rate-down

Context config>subscr-mgmt>loc-user-db>ppp>host>ale

Description This command is applicable to LAC and LNS. It provides the last mile link rate in the downstream direction that is needed for proper shaping and calculating the interleaving delay.

The rate information in the last mile is taken from the following sources in the order of priority:

- Statically provisioned value in local user database (LUDB).

- RADIUS.
- PPPoE tags on LAC or ICRQ message (RFC 5515) /ICCN message (TX Connect Seed) on LNS.

The **no** form of this command reverts to the default.

Parameters *rate* — Specifies last mile link downstream rate in the access loop.

Values 1 to 100000 kb/s

short-sequence-numbers

Syntax [**no**] **short-sequence-numbers**

Context config>subscr-mgmt>ppp-policy>mlppp

Description This command enables a peer request to send short sequence numbers. This command is applicable to LAC and LNS. By default, MLPPPoX negotiates 24bit long sequence numbers. This command allows this to be changed to shorter, 12-bit sequence numbers.

The **no** form of this command reverts to the default.

5.4.2.1.3 PPP/PPPoE Service Commands

ppp

Syntax [**no**] **ppp**

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command configures PPP parameters.

The **no** form of this command reverts all PPP parameters from the PPP context to their defaults.

Default ppp

pppoe

Syntax [**no**] **pppoe**

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command enables the context to configure PPPoE parameters.

The **no** form of this command reverts all PPPoE parameters from the PPPoE context to their defaults.

Default pppoe

anti-spoof

Syntax **anti-spoof** *pppoe-anti-spoofing-type*
no anti-spoof

Context config>service>ies>sub-if>grp-if>pppoe
config>service>vprn>sub-if>grp-if>pppoe

Description This command specifies the type of PPPoE anti-spoof filtering to use.

The **no** form of this command reverts to the default.

Default anti-spoof mac-sid

Parameters *pppoe-anti-spoofing-type* — Specifies the PPPoE anti-spoof filtering.

Values mac-sid, mac-sid-ip

dhcp-client

Syntax **dhcp-client**

Context config>service>vprn>sub-if>grp-if>pppoe
config>service>ies>sub-if>grp-if>pppoe

Description This command enables the context to configure the PPPoE-to-DHCP options.

ccag-use-origin-sap

Syntax [**no**] **ccag-use-origin-sap**

Context config>service>vprn>sub-if>grp-if>pppoe>dhcp-client
config>service>ies>sub-if>grp-if>pppoe>dhcp-client

Description This command enables the original VPLS SAP to be included in the circuit-id option to send to the DHCP server (in case this interface is connected to a VPLS by a CCA MDA).

The **no** form of this command disables the feature.

policy

Syntax **policy** *ppp-policy-name*
no policy

Context config>service>vprn>sub-if>grp-if>pppoe
config>service>vprn>sub-if>grp-if>ppp
config>service>ies>sub-if>grp-if>pppoe
config>service>ies>sub-if>grp-if>ppp

Description This command specifies the PPPoE policy on this interface.

The **no** form of this command reverts to the default.

Default policy “default”

Parameters *ppp-policy-name* — Specifies the PPP policy name up to 32 characters.

python-policy

Syntax **python-policy** [*policy-name*]
no python-policy

Context config>aaa>diam-peer-plcy

Description This command specified the python-policy for Diameter messages received or transmitted on the diameter peers defined in the diameter-peer-policy.

The **no** form of this command reverts to the default.

Parameters *policy-name* — Specifies the name of the Python policy, up to 32 characters.

sap-session-limit

Syntax **sap-session-limit** *sap-session-limit*
no sap-session-limit

Context config>service>ies>sub-if>grp-if>pppoe
config>service>vprn>sub-if>grp-if>pppoe

Description This command specifies the number of PPPoE hosts per SAP allowed for this group-interface.

The **no** form of this command reverts to the default.

Default sap-session-limit 1

Parameters *sap-session-limit* — Specifies the number of PPPoE hosts per SAP allowed.

Note: The operational maximum value may be smaller due to equipped hardware dependencies.

Values 1 to 131071

session-limit

Syntax **session-limit** *session-limit*
no session-limit

Context config>service>ies>sub-if>grp-if>pppoe
config>service>vprn>sub-if>grp-if>pppoe
config>service>ies>sub-if>pppoe
config>service>vprn>sub-if>pppoe
config>service>ies>sub-if>grp-if>ppp
config>service>vprn>sub-if>grp-if>ppp

Description This command specifies the number of PPPoE hosts allowed for this group interface.

The **no** form of this command reverts to the default.

Default session-limit 1

Parameters *session-limit* — Specifies the number of PPPoE hosts allowed.

Note: The operational maximum value may be smaller due to equipped hardware dependencies.

Values 1 to 131071
1 to 262143 (retail subscriber interface)

user-db

Syntax **user-db** *local-user-db-name*
no user-db

Context config>service>ies>sub-if>grp-if>pppoe
config>service>vprn>sub-if>grp-if>pppoe
config>service>ies>sub-if>grp-if>ppp
config>service>vprn>sub-if>grp-if>ppp

Description This command configures the local user database to use for PPP PAP/CHAP authentication

The **no** form of this command reverts to the default.

Parameters *local-user-db-name* — Specifies the local user database name up to 32 characters.

5.4.2.1.4 RADIUS Attribute Commands

acct-authentic

Syntax [**no**] **acct-authentic**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the acct-authentic RADIUS attribute.

The **no** form of this command reverts to the default.

acct-delay-time

Syntax [**no**] **acct-delay-time**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the acct-delay-time RADIUS attribute.

The **no** form of this command reverts to the default.

called-station-id

Syntax [**no**] **called-station-id**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command includes called station id attributes.

The **no** form of this command reverts to the default.

calling-station-id

Syntax **calling-station-id**
calling-station-id {**mac** | **remote-id** | **sap-id** | **sap-string**}
no calling-station-id

Context config>service>ies>if>sap
config>service>ies>sub-if>grp-if>sap
config>service>vpls>sap
config>service>vprn>if>sap
config>service>vprn>sub-if>grp-if>sap
config>subscr-mgmt>auth-plcy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include>include-radius-attribute

Description This command enables the inclusion of the calling-station-id attribute in RADIUS authentication requests and RADIUS accounting messages. The value inserted is set at the SAP level. If no **calling-station-id** value is set at the SAP level, the **calling-station-id** attribute is not sent.

The **no** form of this command reverts to the default.

Parameters **mac** — Specifies that the mac-address is sent.

remote-id — Specifies that the remote-id is sent.

sap-id — Specifies that the SAP ID is sent.

sap-string — Specifies that the value is the inserted value set at the SAP level. If no **calling-station-id** value is set at the SAP level, the **calling-station-id** attribute is not sent.

circuit-id

Syntax [**no**] **circuit-id**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the agent-circuit-id for RADIUS.

The **no** form of this command reverts to the default.

delegated-ipv6-prefix

Syntax [**no**] **delegated-ipv6-prefix**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the delegated-ipv6-prefix RADIUS attribute.

The **no** form of this command disables the generation of the delegated-ipv6-prefix RADIUS attribute.

framed-interface-id

Syntax [**no**] **framed-interface-id**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the framed-interface-id RADIUS attribute.

The **no** form of this command disables the framed-interface-id RADIUS attribute.

framed-ip-addr

Syntax [**no**] **framed-ip-addr**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the inclusion of the framed-ip-addr attribute.

The **no** form of this command disables the inclusion of the framed-ip-addr attribute.

framed-ip-netmask

Syntax [**no**] **framed-ip-netmask**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the inclusion of the framed-ip-netmask attribute.

The **no** form of this command disables the inclusion of the framed-ip-netmask attribute.

framed-ipv6-prefix

Syntax [**no**] **framed-ipv6-prefix**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the framed-ipv6-prefix RADIUS attribute.

The **no** form of this command disables the generation of the framed-ipv6-prefix RADIUS attribute.

ipv6-address

Syntax [**no**] **framed-ipv6-address**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the ipv6-address RADIUS attribute.

The **no** form of this command disables the generation of the ipv6-address RADIUS attribute.

mac-address

Syntax [**no**] **mac-address**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the client MAC address RADIUS attribute.

The **no** form of this command disables the generation of the client MAC address RADIUS attribute.

nas-identifier

Syntax [**no**] **nas-identifier**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the nas-identifier RADIUS attribute.

The **no** form of this command disables the generation of the nas-identifier RADIUS attribute.

nas-port

Syntax [**no**] **nas-port** *bit-specification binary-spec*

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the nas-port RADIUS attribute. You enter decimal representation of a 32-bit string that indicates your port information. This 32-bit string can be compiled based on different information from the port (data types). By using syntax number-of-bits data-type you indicate how many bits from the 32 bits are used for the specific data type. These data types can be combined up to 32 bits in total. In between the different data types 0's and/or 1's as bits can be added.

The **no** form of this command disables the nas-port configuration.

Parameters *bit-specification binary-spec* — Specifies the NAS port attribute.

Values

binary-spec	<bit-specification> <binary-spec>
bit-specification	0 | 1 | <bit-origin>
bit-origin	*<number-of-bits><origin>
number-of-bits	1 to 32
origin	o | i | s | m | p
	o outer VLAN ID
	i inner VLAN ID
	s slot number
	m MDA number
	p port number or lag-id

Output

Sample

```
*12o*12i00*2s*2m*2p => oooo oooo oooo iiii iiii iiii 00ss mmpp
If outer vlan = 0 & inner vlan = 1 & slot = 3 & mda = 1 & port = 1
=>  0000 0000 0000 0000 0000 0001 0011 0101 => nas-port = 309
```

nas-port-id

Syntax [**no**] **nas-port-id** [**prefix-string** *string*] [**suffix** *suffix-option*]

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the nas-port-id RADIUS attribute. Optionally, the value of this attribute (the SAP ID) can be prefixed by a fixed string and suffixed by the circuit-id or the remote-id of the client connection. If a suffix is configured, but no corresponding data is available, the suffix used is 0/0/0/0/0/0.

The **no** form of this command disables the generation of the nas-port-id RADIUS attribute.

Parameters *string* — Specifies that a user configurable string is added to the RADIUS NAS port attribute, up to 8 characters.

suffix-option — Specifies the suffix type to be added to the RADIUS NAS port attribute.

Values circuit-id, remote-id

nas-port-type

Syntax **nas-port-type**
nas-port-type *value*
no nas-port-type

Context config>subscr-mgmt>auth-plcy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the nas-port-type RADIUS attribute. If set to **nas-port-type**, the following is sent: values: 32 (null-encap), 33 (dot1q), 34 (qinq), 15 (DHCP hosts). The **nas-port-type** can also be set as a specified value, with an integer from 0 to 255.

The **no** form of this command disables the generation of the nas-port-type RADIUS attribute

Parameters *value* — Specifies an enumerated integer that specifies the value that is put in the RADIUS nas-port-type attribute.

Values 0 to 255

nat-port-range

Syntax [**no**] **nat-port-range**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the of nat-port-range attribute.

The **no** form of this command disables the generation of the nat-port-range attribute.

remote-id

Syntax [**no**] **remote-id**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the agent-remote-id for RADIUS.

The **no** form of this command disables the generation of the agent-remote-id for RADIUS.

sap-session-index

Syntax [**no**] **sap-session-index**

Context config>subscr-mgmt>auth-policy>include-radius-attribute

Description This command includes sap-session-index attributes.

The **no** form of this command excludes sap-session-index attributes.

sla-profile

Syntax [**no**] **sla-profile**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command specifies that SLA profile attributes should be included into RADIUS accounting messages.

The **no** form of this command excludes SLA profile attributes into RADIUS accounting messages.

sub-profile

Syntax [**no**] **sub-profile**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command specifies that subscriber profile attributes should be included into RADIUS accounting messages.

The **no** form of this command excludes subscriber profile attributes into RADIUS accounting messages.

subscriber-id

Syntax [**no**] **subscriber-id**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command specifies that subscriber ID attributes should be included into RADIUS accounting messages.

The **no** form of this command excludes subscriber ID attributes into RADIUS accounting messages.

radius-accounting-server

Syntax **radius-accounting-server**

Context config>app-assure>rad-acct-plcy
config>aaa>l2tp-tunnel-acct-plcy

Description This command creates the context for defining RADIUS accounting server attributes under a given session authentication policy.

access-algorithm

Syntax **access-algorithm** {**direct** | **round-robin**}

Context config>app-assure>rad-acct-plcy>server

Description This command configures the algorithm used to access the list of configured RADIUS servers.

Default access-algorithm direct

Parameters **direct** — Specifies that the first server is used as primary server for all requests, the second as secondary and so on.

round-robin — Specifies that the first server is used as primary server for the first request, the second server as primary for the second request, and so on. If the router gets to the end of the list, it starts again with the first server.

retry

Syntax **retry** *count*

Context config>app-assure>rad-acct-plcy>server

Description This command configures the number of times the router attempts to contact the RADIUS server for authentication.

The **no** form of this command reverts to the default.

Note: The retry count includes the first attempt.

Default retry 3 (the initial attempt as well as two retried attempts)

Parameters *count* — Specifies the retry count.

Values 1 to 10

router

Syntax **router** *router-instance*
router service-name *service-name*

no router

Context config>app-assure>rad-acct-plcy>server

Description This command specifies the number of times the router attempts to contact the RADIUS server for authentication, if not successful the first time.

The **no** form of this command reverts to the default value.

server

Syntax **server** *server-index* **address** *ip-address* **secret** *key* [**hash** | **hash2**] [**port** *port*] [**create**]
no server *server-index*

Context config>app-assure>rad-acct-plcy>server

Description This command adds a RADIUS server and configures the RADIUS server IP address, index, and key values.

Up to five RADIUS servers can be configured at any one time. RADIUS servers are accessed in order from lowest to highest index for authentication requests until a response from a server is received. A higher indexed server is only queried if no response is received from a lower indexed server (which implies that the server is not available). If a response from a server is received, no other RADIUS servers are queried.

The **no** form of this command removes the server from the configuration.

Parameters *server-index* — The index for the RADIUS server. The index determines the sequence in which the servers are queried for authentication requests. Servers are queried in order from lowest to highest index.

Values 1 to 16 (a maximum of 5 accounting servers)

ip-address — The IP address of the RADIUS server. Two RADIUS servers cannot have the same IP address. An error message is generated if the server address is a duplicate.

secret *key* — The secret key to access the RADIUS server. This secret key must match the password on the RADIUS server.

Values secret-key — A string up to 20 characters
hash-key — A string up to 33 characters
hash2-key — A string up to 55 characters

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

port — Specifies the UDP port number on which to contact the RADIUS server for authentication.

Values 1 to 65535

source-address-range

Syntax **source-address-range** *start-ip-address end-ip-address*
no source-address

Context config>aaa>isa-radius-plcy>servers

Description This command configures the source address range of the RADIUS messages.

The **no** form of this command reverts to the default value.

Parameters *start-ip-address* — Specifies the start of the range of source addresses to be used for NAT RADIUS accounting.

end-ip-address — Specifies the end of the range of source addresses to be used for NAT RADIUS accounting.

timeout

Syntax **timeout** *seconds*

Context config>app-assure>rad-acct-plcy>server

Description This command configures the number of seconds the router waits for a response from a RADIUS server.

The **no** form of this command reverts to the default value.

Default timeout 5

Parameters *seconds* — Specifies the time the router waits for a response from a RADIUS server.

Values 1 to 90

5.4.2.2 Service Commands

subscriber-interface

Syntax **subscriber-interface** *ip-int-name* [**create**]
subscriber-interface *ip-int-name* [**create**] **fwd-service** *service-id* **fwd-subscriber-interface** *ip-int-name*]
no subscriber-interface *ip-int-name*

Context config>service>ies
config>service>vprn

Description This command allows the operator to create special subscriber-based interfaces. It is used to contain multiple group interfaces. Multiple subnets associated with the subscriber interface can be applied to any of the contained group interfaces in any combination. The subscriber interface allows subnet sharing between group interfaces.

The **no** form of this command reverts to the default.

Parameters *ip-int-name* — Specifies the interface name of a subscriber interface. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

service-id — Specifies the wholesale service ID or service name.

Values *service-id*: 1 to 214748364
svc-name: A string up to 64 characters

ip-int-name — Specifies the wholesale subscriber interface.

group-interface

Syntax **group-interface** *ip-int-name* [**create**] [**type**]
group-interface *ip-int-name*

Context config>service>ies>subscriber-interface
config>service>vprn>subscriber-interface

Description This command creates a group interface. This interface is designed for triple-play services where multiple SAPs are part of the same subnet. A group interface may contain one or more SAPs.

The **no** form of this command removes the group interface from the subscriber interface.

Parameters *ip-int-name* — Specifies the interface name of a group interface. If the string contains special characters (#, $, spaces, and so on.), the entire string must be enclosed within double quotes.

bonding — Specifies to use connection bonding.

gtp — Specifies to use GTP.

lns — Specifies to use LNS.

wlangw — Specifies to use dynamic GRE encapsulation.

dhcp

Syntax **dhcp**

Context config>service>ies>subscriber-interface
config>service>vprn>subscriber-interface

Description This command enables the context to configure DHCP parameters.

client-applications

Syntax **client-applications** {[**dhcp**] [**ppp**]}
no client-applications

Context config>service>ies>subscriber-interface
config>service>vprn>subscriber-interface

Description This command enables DHCP relay and proxy-server for the configured client types.

The **no** form of this command reverts to the default.

Default dhcp

Parameters **dhcp** — Enables IPoE clients to use the DHCP relay or proxy-server.

ppp — Enables PPPoE clients to use the DHCP relay or proxy-server that PPPoE will attempt to request an IP address for a PPPoE client from the DHCP server(s) assigned to PPPoE node.

5.4.2.3 Show Commands

Note: Command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

peer

Syntax **peer** *ip-address* [**udp-port** *port*]

peer *ip-address* statistics [**udp-port** *port*]
peer [**draining**] [**blacklisted** | **selectable** | **unreachable**]

Context show>router>l2tp

Description This command displays L2TP peer operational information/

Parameters *ip-address* — Specifies the IP dress for the L2TP peer.

Values

ip-address	ipv4-address - a.b.c.d
ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D
draining	keyword
statistics	keyword
port	[1 to 65535]

port — Specifies the local UDP port of this L2TP

Values 1 to 65535

draining — Specifies to display information about the L2TP peer being drained.

blacklisted — Specifies to display information about the peer has been deemed unreachable and has been put in the tunnel-selection-blacklist.

selectable — Specifies to display information about the peer has been deemed unreachable for the time specified in tmnxL2tpXtTuSelBlacklistMaxTime and can now be selected for a single session, after which the reachability of the peer is reevaluated.

unreachable — Specifies to display information about the sysUpTime when the this peer was deemed unreachable for the last time.

Output The following output is an example of L2TP peer information.

Sample Output

```
show router l2tp peer 10.100.0.2
===============================================================================
Peer IP: 10.100.0.2
===============================================================================
Roles capab/actual: LAC LNS /LAC  -     Draining          : false
Tunnels           : 1                   Tunnels Active    : 0
Sessions          : 1                   Sessions Active   : 0
Reachability      : blacklisted         Time Unreachable  : 01/31/2013 08:55:06
Time Blacklisted  : 01/31/2013 08:55:06 Remaining (s)     : 34
===============================================================================
Conn ID                      Loc-Tu-ID Rem-Tu-ID State              Ses Active
  Group                                                            Ses Total
    Assignment
```

```
-------------------------------------------------------------------------------
977207296                   14911     0         closed             0
  base_lac_base_lns                                              1
    t1
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================


show router l2tp tunnel detail
===============================================================================
L2TP Tunnel Status
===============================================================================
Connection ID: 831782912
State        : closedByPeer
IP           : 10.0.0.1
Peer IP      : 10.100.0.2
Tx dst-IP    : 10.100.0.2
Rx src-IP    : 10.100.0.2
Name         : lac
Remote Name  :
Assignment ID: t1
Group Name   : base_lac_base_lns
Acct. Policy : l2tp-base
Error Message: N/A

                                     Remote Conn ID    : 4294901760
Tunnel ID         : 12692            Remote Tunnel ID  : 65535
UDP Port          : 1701             Remote UDP Port   : 1701
Preference        : 50               Receive Window    : 64
Hello Interval (s): 300
Idle TO (s)       : 5                Destruct TO (s)   : 60
Max Retr Estab    : 5                Max Retr Not Estab: 5
Session Limit     : 32767            AVP Hiding        : sensitive
Transport Type    : udpIp            Challenge         : never
Time Started      : 01/31/2013 08:56:58 Time Idle        : 01/31/2013 08:56:58
Time Established  : N/A              Time Closed       : 01/31/2013 08:56:58
Stop CCN Result   : reqShutDown      General Error     : noError
Blacklist-state   : blacklisted
Blacklist Time    : 01/31/2013 08:56:58 Remaining (s)    : 49
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
```

ppp-policy

Syntax **ppp-policy** [*ppp-policy-name* [**association**]]

Context show>subscr-mgmt

Description This command displays PPP policy information.

Parameters *ppp-policy-name* — Specifies an existing PPP policy.

association — Displays the object the PPP policy is associated.

Output The following output displays PPP policy information.

Sample Output

```
*A:ALA-49>show>subscr-mgmt# ppp-policy policy1
===============================================================================
PPP Policy "policy1"
===============================================================================
Last Mgmt Change     : 11/16/2003 20:06:39     PPP-mtu              : N/A
Keepalive Interval   : 10s                     Keepalive Multiplier : 1
Disable AC-Cookies   : No                      PADO Delay           : 0msec
Max Sessions-Per-Mac : 63                      Reply-On-PADT        : No
-------------------------------------------------------------------------------
PPP Custom Options
-------------------------------------------------------------------------------
Protocol Number Value
-------------------------------------------------------------------------------
No options configured.
===============================================================================
*A:ALA-49>show>subscr-mgmt# pppoe-policy policy1 association
===============================================================================
PPPoE Policy "policy1"
===============================================================================
-------------------------------------------------------------------------------
Interface Associations
-------------------------------------------------------------------------------
Service-Id : 20 (IES)
 - grp_pppoe1
 - grp_pppoe2
 - grp_pppoe3
===============================================================================
*A:ALA-49>show>subscr-mgmt#
```

ipoe

Syntax **ipoe**

Context show>service>id

Description This command enables the context to display IPoE information.

session

Syntax **session** [**sap** *sap-id*] [**mac** *ieee-address*] [**circuit-id** *circuit-id*] [**remote-id** *remote-id*] [**interface** *ip-int-name* | *ip-address*] [**inter-dest-id** *intermediate-destination-id*] [**no-inter-dest-id**] [**ip-address** *ip-prefix*[/*prefix-length*]] [**port** *port-id*] [**subscriber** *sub-ident-string*] [**sap-session-id** *sap-session-index*] [**detail**] [**wholesaler** *service-id*] [**router-advertisement-policy** *policy-name*]

Context show>service>id>ipoe

Description This command displays the identified IPoE session details active on the specified service instance.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

ieee-address — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

circuit-id — Specifies the circuit ID up to 127 characters.

remote-id — Specifies that information that goes into the remote-id sub-option in the DHCP relay packet, up to 255 characters.

ip-int-name — Specifies the name of the IP interface up to 32 characters.

ip-address[/prefix-length] — Specifies information for the specified IP address and mask.

no-inter-dest-id — Displays the information about no intermediate destination ID.

port-id — Displays information about the specified port ID.

sap-session-index — Displays sap-session-index information.

service-id — Specifies the service ID of the wholesaler. This parameter applies only to the 7750 SR.

Values

service-id:	1 to 2147483647

intermediate-destination-id — Specifies the intermediate destination identifier which is encoded in the identification strings up to 32 characters.

sub-ident-string — Specifies an existing subscriber identification profile.

detail — Displays all IPoE session details.

policy-name — Specifies the name of the router advertisement policy, up to 32 characters.

Output The following is an example of service IPoE session information.

Sample Output

```
# show service id 4000 ipoe session
===============================================================================
IPoE sessions for svc-id 4000
===============================================================================
Sap Id                          Mac Address         Up Time         MC-Stdby
    Subscriber-Id
        [CircuitID] | [RemoteID]
-------------------------------------------------------------------------------
1/1/4:1201.27                   00:51:00:00:00:0c   0d 00:00:18
    ipoe-session-001
-------------------------------------------------------------------------------
CID | RID displayed when included in session-key
```

```
Number of sessions : 1
==============================================================================


# show service id 4000 ipoe session detail
===============================================================================
IPoE sessions for service 4000
===============================================================================
SAP                     : 1/1/4:1201.27
Mac Address             : 00:51:00:00:00:0c
Circuit-Id              : circuit-id-1
Remote-Id               : remote-id-1
Session Key             : sap-mac

MC-Standby              : No

Subscriber-interface    : sub-int-1
Group-interface         : group-int-1

Up Time                 : 0d 00:01:01
Session Time Left       : N/A
Last Auth Time          : 02/28/2015 01:01:09
Min Auth Intvl (left)   : 0d 00:05:00 (0d 00:03:59)
Persistence Key         : N/A

Subscriber              : "ipoe-session-001"
Sub-Profile-String      : "sub-profile-1"
SLA-Profile-String      : "sla-profile-1"
ANCP-String             : ""
Int-Dest-Id             : ""
App-Profile-String      : ""
Category-Map-Name       : ""
Acct-Session-Id         : "144DFF0000001354D806D5"
Sap-Session-Index       : 1

IP Address              : 10.10.1.201/24
IP Origin               : Radius
Primary DNS             : N/A
Secondary DNS           : N/A
Primary NBNS            : N/A
Secondary NBNS          : N/A
Address-Pool            : N/A

IPv6 Prefix             : 2001:db8:a:111::/64
IPv6 Prefix Origin      : Radius
IPv6 Prefix Pool        : ""
IPv6 Del.Pfx.           : 2001:db8:a001:a100::/56
IPv6 Del.Pfx. Origin    : Radius
IPv6 Del.Pfx. Pool      : ""
IPv6 Address            : 2001:db8:a:101::aaa:1
IPv6 Address Origin     : Radius
IPv6 Address Pool       : ""
Primary IPv6 DNS        : N/A
Secondary IPv6 DNS      : N/A

Radius Session-TO       : N/A
Radius Class            :
Radius User-Name        : 00:51:00:00:00:0c
-------------------------------------------------------------------------------
```

```
Number of sessions : 1
==============================================================================


*A:Dut-A# show service id 13 ipoe session router-advertisement-policy ra-policy-01
==============================================================================
IPoE sessions for svc-id 13
==============================================================================
Sap Id                          Mac Address         Up Time         MC-Stdby
    Subscriber-Id
        [CircuitID] | [RemoteID]
------------------------------------------------------------------------------
1/1/1:13                        00:00:00:00:00:13   0d 00:13:06
    sub-1
1/1/1:13                        00:00:00:00:00:14   0d 00:08:03
    sub-1
1/1/1:13                        00:00:00:00:00:15   0d 00:06:23
    sub-1
------------------------------------------------------------------------------
CID | RID displayed when included in session-key
Number of sessions : 3
==============================================================================
*A:Dut-A#
```

summary

Syntax **summary**

Context show>service>id>ipoe

Description This command displays a summary of IPoE information.

pppoe

Syntax **pppoe**

Context show>service>id

Description This command enables the context to display PPPoE information.

session

Syntax **session** [**interface** *ip-int-name* | *ip-address* | **sap** *sap-id*] [**session-id** *session-id*] [**mac** *ieee-address*] [**ip-address** *ip-address*[/*mask*]] [**port** *port-id*] [**no-inter-dest-id** | **inter-dest-id** *intermediate-destination-id*] [**steering-profile** *steering-profile*] [**router-advertisement-policy** *policy-name*] [**detail** | **statistics**]

session *l2tp-connection-id connection-id* [**detail** | **statistics**]

Context show>service>id>pppoe

Description This command displays PPPoE session information.

Parameters *ip-int-name* — Specifies the IP interface name.

ip-address — Specifies the IP address of the PPPoE session.

sap-id — Specifies the SAP ID.

session-id — Specifies the ID of the PPPoE session.

ieee-address — Specifies the MAC address of the PPPoE session.

port-id — Specifies the port ID.

no-inter-dest-id — Displays the information about no intermediate destination ID.

intermediate-destination-id — Specifies the intermediate destination ID.

steering-profile — Specifies the name of the steering profile, up to a maximum of 32 characters.

policy-name — Specifies the name of the router advertisement policy, up to a maximum of 32 characters.

detail — Displays detailed information.

statistics — Displays statistics about the PPPoE session.

Output The following output displays PPPoE session information.

Sample Output

```
*A:ALA-49#show service id 20 pppoe session
===============================================================================
PPPoE sessions for svc-id 20
===============================================================================
Sap Id              Mac Address       Sid Up Time        IP Address
-------------------------------------------------------------------------------
1/1/3:200           00:00:00:00:00:03 1   1d 00:48:39    10.0.0.101
1/1/3:300           00:00:00:00:00:05 1   0d 00:01:08    10.0.0.119
-------------------------------------------------------------------------------
Number of sessions : 2
===============================================================================
*A:ALA-49#


*A:ALA-49# show service id 20 pppoe session ip-address 20.0.0.101 detail
===============================================================================
PPPoE sessions for svc-id 20
===============================================================================
Sap Id              Mac Address       Sid Up Time        IP Address
-------------------------------------------------------------------------------
1/1/3:200           00:00:00:00:00:03 1   1d 00:49:46    10.0.0.101

LCP State            : Opened
IPCP State           : Opened
PPP MTU              : 1492
PPP Auth-Protocol    : PAP
PPP User-Name        : user4@domain1
```

```
Subscriber-interface : sub_pppoe
Group-interface      : grp_pppoe2

Subscriber Origin    : RADIUS
Strings Origin       : RADIUS
IPCP Info Origin     : DHCP

Subscriber           : "radius_papchap4"
Sub-Profile-String   : "sub1"
SLA-Profile-String   : "sla1"
ANCP-String          : ""
Int-Dest-Id          : ""
App-Profile-String   : ""

Primary DNS          : N/A
Secondary DNS        : N/A
Primary NBNS         : N/A
Secondary NBNS       : N/A

Circuit-Id           : 2
Remote-Id            :

Session-Timeout      : N/A
-------------------------------------------------------------------------------
Number of sessions : 1
===============================================================================
*A:ALA-49#


*A:ALA-49# show service id 20 pppoe session ip-address 20.0.0.101 statistics
===============================================================================
PPPoE sessions for svc-id 20
===============================================================================
Sap Id              Mac Address       Sid Up Time        IP Address
-------------------------------------------------------------------------------
1/1/3:200           00:00:00:00:00:03 1   1d 00:50:39    10.0.0.101

Packet Type               Received        Transmitted
-------------------------------------------------------------------------------
LCP Configure-Request     1               2
LCP Configure-Ack         1               1
LCP Configure-Nak         1               0
LCP Configure-Reject      0               0
LCP Terminate-Request     0               0
LCP Terminate-Ack         0               0
LCP Code-Reject           0               0
LCP Echo-Request          8927            866
LCP Echo-Reply            866             8927
LCP Protocol-Reject       0               0
LCP Discard-Request       0               0
-------------------------------------------------------------------------------
PAP Authenticate-Request  1               -
PAP Authenticate-Ack      -               1
PAP Authenticate-Nak      -               0
-------------------------------------------------------------------------------
CHAP Challenge            -               0
CHAP Response             0               -
CHAP Success              -               0
CHAP Failure              -               0
```

```
-------------------------------------------------------------------------------
IPCP Configure-Request    2               1
IPCP Configure-Ack        1               1
IPCP Configure-Nak        0               1
IPCP Configure-Reject     0               0
IPCP Terminate-Request    0               0
IPCP Terminate-Ack        0               0
IPCP Code-Reject          0               0
-------------------------------------------------------------------------------
Unknown Protocol          0               -
-------------------------------------------------------------------------------
Number of sessions : 1
===============================================================================
*A:ALA-49#


*A:Dut-C# show service id 2000 pppoe session detail
===============================================================================
PPPoE sessions for svc-id 2000
===============================================================================
Sap Id              Mac Address       Sid    Up Time               Type
    IP/L2TP-Id/Interface-Id
-------------------------------------------------------------------------------
2/1/5:2000          00:01:00:00:04:15 1      0d 00:05:07           Local
    239.1.5.22

LCP State            : Opened
IPCP State           : Opened
IPv6CP State         : Initial
PPP MTU              : 1492
PPP Auth-Protocol    : None
PPP User-Name        : (Not Specified)

Subscriber-interface : ies-2000-239.1.1.1
Group-interface      : grp-Vprn-2/1/5

Subscriber Origin    : RADIUS
Strings Origin       : RADIUS
IPCP Info Origin     : RADIUS
IPv6CP Info Origin   : None

Subscriber           : "hpolSub43"
Sub-Profile-String   : "hpolSubProf2"
SLA-Profile-String   : "hpolSlaProf1"
ANCP-String          : ""
Int-Dest-Id          : "2000"
App-Profile-String   : ""
Category-Map-Name    : ""

Primary DNS          : N/A
Secondary DNS        : N/A
Primary NBNS         : N/A
Secondary NBNS       : N/A
Address-Pool         : N/A

IPv6 Prefix          : N/A
IPv6 Del.Pfx.        : N/A
Primary IPv6 DNS     : N/A
Secondary IPv6 DNS   : N/A
```

```
Circuit-Id           : circuit 0
Remote-Id            : remote 00-00-00-00-00-00-eth0-2
Service-Name         :

Session-Timeout      : N/A
RADIUS Class         :
RADIUS User-Name     : 00:01:00:00:04:15
Data link            : aal5
Encaps 1             : notAvailable
Encaps 2             : pppoaLlc
-------------------------------------------------------------------------------
Overrides
-------------------------------------------------------------------------------
Direction Type           Key        PIR        CIR        CBS        MBS
-------------------------------------------------------------------------------
Egress    Agg-Rate-Limit N/A        24125940   N/A        N/A        N/A
-------------------------------------------------------------------------------
No. of Overrides: 1
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Number of sessions   : 1
===============================================================================
*A:Dut-C#

*A:Dut-A#show  service id 10 pppoe session router-advertisement-policy "ra-policy-
01"
===============================================================================
PPPoE sessions for svc-id 1000
===============================================================================
Sap Id              Mac Address       Sid    Up Time         Type
    IP/L2TP-Id/Interface-Id                                    MC-Stdby
-------------------------------------------------------------------------------
1/1/20:841          00:00:64:19:01:02 1      0d 00:00:18     local
    192.168.0.10
    02:00:64:FF:FE:19:01:02
-------------------------------------------------------------------------------
Number of sessions   : 1
===============================================================================
```

statistics

Syntax **statistics** [{**sap** *sap-id* | **interface** *ip-int-name* | *ip-address*}]

Context show>service>id>pppoe

Description This command displays PPPoE statistics.

Parameters *sap-id* — Displays information for the specified SAP.

ip-int-name — Displays information about the specified interface.

ip-address — Displays information about the specified IP address.

Output The following output displays PPPoE statistics information.

Sample Output

```
*A:ALA-49# show service id 20 pppoe statistics
===============================================================================
PPPoE statistics for IES service 20
===============================================================================
Packet Type         Received        Transmitted
-------------------------------------------------------------------------------
PADI                2               -
PADO                -               2
PADR                2               -
PADS                -               2
PADT                0               0
session             9838            9839
-------------------------------------------------------------------------------
Drop Counters
-------------------------------------------------------------------------------
Rx Invalid Version   : 0
Rx Invalid Type      : 0
Rx Invalid Code      : 0
Rx Invalid Session   : 0
Rx Invalid Length    : 0
Rx Invalid Tags      : 0
Rx Invalid AC-Cookie : 0
Rx Dropped           : 0
===============================================================================
*A:ALA-49#
```

summary

Syntax **summary**

Context show>service>id>pppoe

Description This command displays PPPoE summary information.

5.4.2.4 Clear Commands

pppoe

Syntax **pppoe**

Context clear>service>id

Description This command enables the context to clear PPPoE-related data for the specified service.

session

Syntax **session all** [**no-padt**]
session {**interface** *ip-int-name* | *ip-address* | **sap** *sap-id*} [**mac** *ieee-address*] [**session-id** *session-id*] [**ip-address** *ip-address*[/*mask*]] [**port** *port-id*] [**no-inter-dest-id** | **inter-dest-id** *intermediate-destination-id*] [**no-padt**]

Context clear>service>id>ppoe

Description This command clears PPPoE sessions.

Parameters *ip-int-name* — Clears session information for the specified IP interface name.

ip-address — Clears session information for the specified IP interface name.

mask — Clears session information for the specified mask.

sap-id — Clears session information for the specified SAP ID.

ieee-address — Clears session information for the specified MAC address.

session-id — Clears session information for the specified session ID.

port-id — Clears session information for the specified port ID.

no-inter-dest-id — Clears session information for the specified intermediate destination ID.

inter-dest-id — Clears session information for the specified intermediate destination ID.

no-padt — Clears session information for the PPPoE Active Discovery Terminate (PADT) packets transmitted by this PPPoE client.

statistics

Syntax **statistics** [{**sap** *sap-id* | **interface** *ip-int-name* | *ip-address*}]

Context clear>service>id>ppoe

Description This command clears PPPoE statistics.

5.4.2.5 Debug Commands

ppp

Syntax [**no**] **ppp**

Context debug>service>id

Description This command enables and configures PPP debugging.

event

Syntax [**no**] **event**

Context debug>service>id>ppp

Description This command enables debugging for specific PPPoE events.

dhcp-client

Syntax **dhcp-client** [**terminate-only**]
no dhcp-client

Context debug>service>id>ppp>event

Description This command enables debugging for specific DHCP client events.

Parameters **terminate-only** — Displays terminate-only DHCP client information.

ppp

Syntax **ppp** [**terminate-only**]
no ppp

Context debug>service>id>ppp>event

Description This command enables debugging for PPP events.

Parameters **terminate-only** — Enables debugging for terminate-only PPP events.

mac

Syntax [**no**] **mac** *ieee-address*

Context debug>service>id>ppp

Description This command shows PPP packets for a particular MAC address.

Parameters *ieee-address* — Enables debugging for the specified MAC address.

packet

Syntax [**no**] **packet**

Context debug>service>id>ppp

Description This command enables debugging for specific PPPoE packets.

detail-level

Syntax **detail-level** {**low** | **medium** | **high**}
no detail-level

Context debug>service>id>ppp>packet

Description This command configures the PPP packet tracing detail level.

Parameters **low, medium, high** — Enables debugging for the PPP packet tracing detail level.

dhcp-client

Syntax [**no**] **dhcp-client**

Context debug>service>id>ppp>packet

Description This command enables debugging for specific DHCP client packets.

discovery

Syntax **discovery** [**padi**] [**pado**] [**padr**] [**pads**] [**padt**]
no discovery

Context debug>service>id>ppp>packet

Description This command enables debugging for specific PPP discovery packets.

Parameters **padi** — Enables debugging for PADI PPP discovery packets.

pado — Enables debugging for PADO PPP discovery packets.

padr — Enables debugging for PADR PPP discovery packets.

pads — Enables debugging for PADS PPP discovery packets.

padt — Enables debugging for PADT PPP discovery packets.

mode

Syntax **mode** {**dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}
no mode

Context debug>service>id>ppp>packet

Description This command configures the PPP packet tracing mode.

Parameters **dropped-only** — Enables debugging for dropped packets.

ingr-and-dropped — Enables debugging for ingress and dropped packets.

egr-ingr-and-dropped — Enables debugging for egress, ingress and dropped packets.

ppp

Syntax **ppp** [**lcp**] [**pap**] [**chap**] [**ipcp**]
no ppp

Context debug>service>id>ppp>packet

Description This command enables debugging for specific PPP packets

Parameters **lcp** — Enables debugging for LCP packets.

pap — Enables debugging for PAP packets.

chap — Enables debugging for CHAP packets.

ipcp — Enables debugging for IPCP packets.

sap

Syntax [**no**] **sap** *sap-id*

Context debug>service>id>ppp

Description This command debugs PPP packets for a particular SAP.

Parameters *sap-id* — Enables debugging for the specified SAP ID.

6 Layer 2 Tunneling Protocol (L2TP)

This chapter provides information about using L2TP, including theory, supported features and configuration process overview.

Note: The information in this section applies only to the 7750 SR.

6.1 Terminology

- Tunnel spec — Describes the requirements for a tunnel and is defined as a set of parameters that will be used in tunnel setup/selection process. The tunnel-spec is defined in the CLI or can be supplied through RADIUS.
- Tunnel (instance) — A run-time object with a unique id terminating at a specific peer. Any change in the tunnel spec once the tunnel has been created has no bearings on the tunnel itself. The list of tunnels can be obtained using the **show router l2tp tunnel** command.
- Peer — A run-time object that is defined by a *ip-address/port* combination. Multiple tunnels can be terminated on the same peer. The list of peers can be obtained using the **show router l2tp peer** command.

6.2 L2TP Overview

6.2.1 LAC DF Bit

The Layer 2 access concentrator (LAC) DF bit is configurable, but by default, it sends all L2TP packets with the DF bit set to 1. Clearing the DF bit will allow downstream routers to fragment the L2TP packets. The LAC itself will not fragment L2TP packets. L2TP packets that have a larger MTU size than what the LAC egress ports allows are dropped. The DF bit can also be configured via RADIUS attribute **Alc-Tunnel-DF-bit**.

6.2.2 Handling L2TP Tunnel/Session Initialization Failures

6.2.2.1 L2TP Tunnel/Session Initialization Failover Mechanisms on LAC

In deployment scenarios with multiple LNS nodes, a list of those LNS nodes can be presented to the LAC during the L2TP session instantiation process (either through CLI or RADIUS). An example of this would be a RADIUS Accept message with a list of tunnel peers:

```
tunnel.com Auth-Type := Local, Password == "tunnel1"
Tunnel-Type:1 += L2TP,
       Tunnel-Medium-Type:1 += IP,
       Tunnel-Client-Auth-Id:1 += lns_tun,
       Tunnel-Assignment-Id:1 += 1,
     Tunnel-Client-Endpoint:1 += 10.0.0.1,
     Tunnel-Server-Endpoint:1 += 10.0.0.2,
       Tunnel-Password:1 += TUNNELPASS,

          Tunnel-Type:2 += L2TP,
     Tunnel-Medium-Type:2 += IP,
         Tunnel-Client-Auth-Id:2 += lns_tun,
     Tunnel-Assignment-Id:2 += 2,
        Tunnel-Client-Endpoint:2 += 10.0.0.1,
        Tunnel-Server-Endpoint:2 += 10.0.0.3,
        Tunnel-Password:2 += TUNNELPASS,

               Tunnel-Type:3 += L2TP,
        Tunnel-Medium-Type:3 += IP,
        Tunnel-Client-Auth-Id:3 += lns_tun,
        Tunnel-Assignment-Id:3 += 3,
        Tunnel-Client-Endpoint:3 += 10.0.0.1,
        Tunnel-Server-Endpoint:3 += 10.0.0.4,
        Tunnel-Password:3 += TUNNELPASS,
               Tunnel-Type:4 += L2TP,

        Tunnel-Medium-Type:4 += IP,
        Tunnel-Client-Auth-Id:4 += lns_tun,
        Tunnel-Assignment-Id:4 += 4,
        Tunnel-Client-Endpoint:4 += 10.0.0.1,
        Tunnel-Server-Endpoint:4 += 10.0.0.5,
        Tunnel-Password:4 += TUNNELPASS
```

If the tunnel or the session establishment attempt fails for any reason, a search for additional operational facilities (tunnels or peers) is made in order to complete the establishment of the tunnel or session that failed in the previous attempt. Moreover, sometimes it is required to go beyond this automatic search for the new facilities and place the tunnel/peer in question into a blacklist. A tunnel timeout always forces the corresponding peer and the tunnel into the blacklist. In addition, a tunnel can be forced into the blacklist by certain explicit error codes (CDN, and Stop-CCN) during the tunnel/session initialization phase. A peer is never forced on a blacklist as a consequence of explicit Result-Code sent by LNS.

Blacklisted peers and tunnels are not eligible to serve new incoming L2TP session until they are removed from the blacklist. The exception is when all tunnel specs evaluate into a blacklisted item. Then, a blacklisted item (tunnel) is tried.

6.2.2.2 Peer Blacklist

A peer is always placed into the blacklist if:

- An attempt to establish a new tunnel fails due to a time out (SCCRQ and SCCCN timeouts)
- The timeout occurs on any control packet within an already established tunnel. All sessions on such tunnel are terminated (PADT is sent toward the clients, StopCCN is sent toward the LNS). Other tunnels that are terminated on the same peer will timeout on their own (if the peer is indeed non-operational), for example, the 7750 SR will not explicitly tear them down based on the timeout of a single tunnel. The timeout of an existing tunnel is caused by lack of acknowledgments to transmitted control packets (ICRQ, ICCN, CDN, Hello).

A tunnel timeout will occur if an acknowledgment is not received after max-retries-established (on an established tunnel) or max-retries-not-established (for the tunnel in the process of being established) retries.

Although there is no configuration option that would control whether a peer can or cannot be blacklisted (it is always blacklisted on tunnel timeout), the amount of time that a peer remains in the blacklist is configurable within the **tunnel-selection-blacklist** CLI node.

6.2.2.3 Tunnel Blacklists

A tunnel spec (that evaluates into a tunnel) is temporary unusable in case that corresponding peer or the tunnel is blacklisted. The following events will trigger placement of the tunnel into the blacklist:

1. Explicit termination of the L2TP session that is in the process of being established within this tunnel. The following CDN Result Codes will place a tunnel to a blacklist (text in red are CLI keywords that will enable specific Result Codes as triggers and [rx,tx] is direction of the messages from the LAC perspective):
   - 02 DisconnectedSeeErrorCode, rx (cdn-err-code)
   - 04 TempMissingFacilities, rx (cdn-tmp-no-facilities)

     Transmit CDN when no session can be allocated

     Audit not yet complete
   - 05 PermanentMissingFacilities, rx (cdn-perm-no-facilities)

     No result code available
   - 06 InvalidDestination, rx(cdn-inv-dest)

     Tunnel is not usable (for example lns-group is not configure on LNS)
   - 10 NotEstablishedInAllotedTime, tx (tx-cdn-not-established-in-time)
2. Explicit termination of the L2TP tunnel in the process of establishment via Stop-CCN Result-Codes:
   - (1) General request to clear control connection, rx (stop-ccn-other)
   - (2) General error, rx (stop-ccn-err-code)
   - (4) Requester is not authorized to establish a control channel, rx, tx (stop-ccn-other)
   - (5) Protocol version not supported, rx, tx (stop-ccn-other)
   - (6) Requester is being shutdown, rx (stop-ccn-other)

Error messages identified by the received Result-Codes can be interpreted as the inability of the LNS to accept additional L2TP sessions within the tunnel (for example due to resource depletion) or to accept additional new tunnels.

The following statements further describe behavior related to the placement of tunnels into the blacklist:

- A new L2TP session establishment attempt will not be triggered on the tunnel that is in the blacklist. Instead, another tunnel will be searched according to the configured preference model.
- The tunnel/session initialization failure will always trigger the selection mechanism for another tunnel. However, it is possible to control via configuration whether to blacklist or not the tunnel for which the L2TP initialization process failed due to certain Result Codes in CDN and/or Stop CCN messages.

- Once the L2TP tunnel/session is established, no events other than the timeout can force the tunnel (and the peer) into the blacklist. In other words, a tunnel Stop or Call disconnect message for a stable tunnel/session will not force the tunnel into the blacklist.
- Existing sessions within the L2TP tunnel will not be purposefully terminated in case that the tunnel is forced into the blacklist due to an explicit reply from LNS indicating the tunnel/session initialization failure. In other words, although the L2TP tunnel might be blacklisted and therefore prevented from serving new L2TP sessions, the existing L2TP session over this tunnel will not be affected.
- A peer will not be forced into the blacklist in case of the explicit failure response from that particular peer. Only tunnels will be blacklisted in that case, assuming that the configuration trigger is enabled. Peers are blacklisted only based on timeouts and not explicit responses.

In case that the end-point is not in the routing table (unreachable via routing), the end-point is marked as permanently unavailable (removed from the L2TP process). Such end-point will never be blacklisted.

6.2.2.3.1 Tunnel Timeout Due to the Peer IP Address Change

In case that the peer address is changed mid-session (for example, from configured IP@ 1.1.1.1 to the new IP@ 2.2.2.2), and then subsequently the tunnel times-out, the new peer 2.2.2.2 would be placed in the blacklist by default. The tunnel itself would not be placed in the blacklist since it is originally tied to a different peer address that it is not in the blacklist. As such it would be eligible for selection the next time a new session request for it arrives. To block selection of this failed tunnel, optionally (by configuration) force it into the blacklist.

This behavior can be enabled with the following CLI:

```
configure router l2tp
configure service vprn <id> l2tp
    tunnel-selection-blacklist
        add-tunnel on <reason> [<reason>...(upto 7 max)]
<reason> : cdn-err-code|cdn-inv-dest|cdn-tmp-no-facilities|cdn-perm-no-
facilities|tx-cdn-not-established-in-time|stop-ccn-err-code|stop-ccn-other|addr-
change-timeout
```

6.2.2.4 Tunnel Selection Mechanism

Once the L2TP tunnel failover is triggered (timeout or specific L2TP session/tunnel setup error message), a new tunnel spec in the list of available tunnel specs will be selected. This tunnel selection mechanism can be controlled via CLI so that the new tunnel-spec is selected from the next preference level. Alternatively the tunnel selection mechanism can be set to a mode where once all the possibilities within the same preference are exhausted, tunnel specs on a higher preference level will be tried.

```
configure router l2tp
configure service vprn <id> l2tp
    next-attempt same-preference-level | next-preference-level
```

In case that all tunnels on a given preference levels are blacklisted, then the behavior will depend on the configuration option as per the following:

- next-attempt = next-preference — Only one tunnel spec from the current preference level will be tried before switching to the next preference level.
- next-attempt = same-preference — All tunnel specs will be tried before switching to the next preference level.

6.2.2.5 Tunnel Probing

Tunnel probing refers to the mechanism where the blacklisted tunnel or an end-point can be selected to serve only a single L2TP session initialization request. Only in case that this single L2TP session is successfully established over the selected tunnel, the tunnel can be removed from the blacklist and consequently can serve new L2TP sessions. The tunnel is eligible for probing once its preconfigured time in the blacklist has expired.

This behavior will ensure that the new session initialization requests are not buffered while waiting for the tunnel to transition into operational state. Buffering would incur session setup delay and in the worst case it would cause session timeout in case that the L2TP tunnel cannot be established.

Without tunnel probing enabled, tunnels will be automatically removed from the blacklist upon the expiry of the preconfigured timer. New consecutive L2TP session initialization requests for such tunnels will always be buffered.

6.2.2.6 Controlling the Size of Blacklist

The size of the blacklist and the time that an item remains ineligible for selection within the blacklist, is configurable.

```
configure router l2tp
configure service vprn <id> l2tp tunnel-selection-blacklist
        max-time    1..60 (minutes)
        max-list-length unlimited | 1..65535
```

6.2.2.7 Displaying the Content of a Blacklist

The content of a blacklist along with the remaining time that each entity is confined to the blacklist can be displayed with the following command:

```
show router <id> l2tp peer blacklisted|not-blacklisted|selectable
```

The following displays blacklist information.

```
show router l2tp peer 10.100.0.2
===============================================================================
Peer IP: 10.100.0.2
===============================================================================
Roles capab/actual: LAC LNS /LAC  -     Draining          : false
Tunnels           : 1                   Tunnels Active    : 0
Sessions          : 1                   Sessions Active   : 0
Reachability      : blacklisted         Time Unreachable  : 01/31/2013 08:55:06
Time Blacklisted  : 01/31/2013 08:55:06 Remaining (s)     : 34
===============================================================================
Conn ID                     Loc-Tu-ID Rem-Tu-ID State              Ses Active
  Group                                                           Ses Total
    Assignment
-------------------------------------------------------------------------------
977207296                   14911     0         closed             0
  base_lac_base_lns                                               1
    t1
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================


show router l2tp tunnel detail
===============================================================================
L2TP Tunnel Status
===============================================================================
Connection ID: 831782912
State        : closedByPeer
IP           : 10.0.0.1
Peer IP      : 10.100.0.2
Tx dst-IP    : 10.100.0.2
Rx src-IP    : 10.100.0.2
```

```
Name         : lac
Remote Name  :
Assignment ID: t1
Group Name   : base_lac_base_lns
Acct. Policy : l2tp-base
Error Message: N/A
                                        Remote Conn ID    : 4294901760
Tunnel ID         : 12692               Remote Tunnel ID  : 65535
UDP Port          : 1701                Remote UDP Port   : 1701
Preference        : 50                  Receive Window    : 64
Hello Interval (s): 300
Idle TO (s)       : 5                   Destruct TO (s)   : 60
Max Retr Estab    : 5                   Max Retr Not Estab: 5
Session Limit     : 32767               AVP Hiding        : sensitive
Transport Type    : udpIp               Challenge         : never
Time Started      : 01/31/2013 08:56:58 Time Idle         : 01/31/2013 08:56:58
Time Established  : N/A                 Time Closed       : 01/31/2013 08:56:58
Stop CCN Result   : reqShutDown         General Error     : noError
Blacklist-state   : blacklisted
Blacklist Time    : 01/31/2013 08:56:58 Remaining (s)     : 49
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
```

6.2.2.8 Generating Trap when the Blacklist is Full

A log is generated when the blacklist reaches its max limit of items. The log event is tmnxL2tpTunnelSelectionBlacklistFull.

6.2.2.9 Premature Removal of Blacklisted Entries

In case that the total number of supported tunnels and peers in blacklist and in the LAC in general has reached its maximum, then on the new session initialization request, the oldest tunnel entry in the blacklist will be removed from the blacklist regardless if the blacklist max-time has expired.

6.2.2.10 Manual Purging of Entities within the Blacklist

The items can be manually purged from the blacklist using the following commands.

```
clear router <id> l2tp tunnel-selection-blacklist
clear router <id> l2tp peer <ip-address> [udp-port <port>] tunnel-selection-
blacklist
clear router <id> l2tp group <tunnel-group-name> [tunnel <tunnel-name>] tunnel-
selection-blacklist
clear router <id> l2tp tunnel <connection-id> tunnel-selection-blacklist
```

6.2.3 CDN Result Code Overwrite

When the number of L2TP sessions reaches the configured maximum value, the LNS sends an out-of-resource Result Code (4 or 5) in a CDN (Call-Disconnect-Notify) message to the LAC. This would trigger the LAC to fail over to another LNS that has the resources available. Similarly, when the tunnel is not usable due to the invalid destination CDN error, the Result-Code 6 will be sent from the LNS.

Certain third-party LAC implementations will trigger tunnel failover only when they receive Result Code 2 in CDN messages (and not 4,5 or 6). In order to support those scenarios, the LNS in the 7750 SR can overwrite result codes 4, 5 and 6 with result code 2 just before they are sent to the LAC. Result Codes can be overwritten only during the L2TP session initialization phase. These codes have the following meanings and are described in RFC 2661, 4.4.2:

- 2 — Call disconnected for the reason indicated in error code
- 4 — Call failed due to lack of appropriate facilities being available (temporary condition)
- 5 — Call failed due to lack of appropriate facilities being available (permanent condition)
- 6 — Invalid Destination

This functionality will be enabled on LNS via the following CLI hierarchy:

```
configure router l2tp
configure service vprn <id> l2tp
   replace-result-code {cdn-tmp-no-facilities | cdn-prem-no-facilities | cdn-inv-
dest}
    no replace-result-code
```

6.2.4 LNS Proxy

LNS offers a proxy LCP (via the **proxy-lcp** command) function where LCP-related information is cached temporarily in the LNS during the ICCN phase where L2TP control messages are exchanged between the LAC and LNS. The LNS can then use the cached information to bypass the LCP negotiation and immediately start the authentication state with the client. Furthermore, proxy authentication (via the **proxy-authentication** command) can also be enabled on the LNS to bypass authentication and the client can immediately start the IPCP negotiation phase.

If the proxy LCP information conflicts from the LNS configuration, then the LNS will force the client to re-start LCP negotiation. LCP negotiation will not be restarted in the proxy LCP mode when:

- MRU is missing in the LastTxLcpConfReq AVP
- The magic number is missing in LastTxLcpConfReq AVP
- Async-Control-Character-Map (ACCM) with value = 0x00000000 is present in LastTxLcpConfReq and LastRxLcpConfReq AVP's
- Address-and-Control-Field-Compression (ACFC) is present in LastTxLcpConfReq and LastRxLcpConfReq. Note that PPP frames with and without address and control field (0xFF03) in the PPP header are accepted. LCP frames without 0xFF03 are also accepted as valid frames.

Also similarly, proxy-authentication that fails will then force the client to re-start the LCP negotiation again.

6.3 L2TP LAC VPRN

Layer 2 Tunneling Protocol (L2TP) allows for PPP sessions to be carried over an IP network.

Each L2TP session transports PPP frames, irrespective of link-layer encapsulation, allows the LNS to terminate PPP sessions that were either PPPoE or PPPoA. L2TP is carried over IPv4 packets in UDP datagrams (default port 1701).

If session data is not reliably delivered, that is, if there is a packet loss, there is no retransmission, a sequence numbers is used within each L2TP session to identify packet loss and re-ordering.

L2TP is comprised of the following concepts:

- L2TP tunnels — L2TP tunnel is a connection between one LAC (L2TP Access Concentrator) and one LNS (L2TP Network Serve) that share a common control channel.
- L2TP sessions — Within each L2TP tunnel, there exists one or more L2TP sessions (one PPP session corresponds to exactly one L2TP session)

L2TP tunnels provide an IP transport for PPP frames between LAC and LNS. In some existing networks, BGP/MPLS VPNs (VPRN in SR-OS) are used to contain the L2TP traffic (and the routes associated with the LAC and LNS) into a dedicated routing instance.

Similar to the LNS implementation, L2TP LAC in a VPRN allows L2TP control and data traffic to be sourced from and received by any valid IP interface within the VPRN (including loopback and interface addresses). L2TP frames may ingress a network port (with up to five MPLS tags) or access ports with SAPs associated with the VPRN IP interfaces.

Non-hitless multi-chassis LAC resiliency

In dual-homed PPPoEv4/v6 wholesale/retail environment over L2TP, the subscriber-hosts are synchronized via Multi-Chassis Synchronization (MCS) protocol. The failover detection mechanism might be implemented via SRRP or Layer 3 MC-LAG with SRRP. When an interface or an entire node fails, the newly selected Master sends PADT to all sessions that were moved over from the failed node.

In case of interface-only failure, CDN is sent towards the LNS to terminate sessions on the LNS.

The PPPoE sessions will be reestablished on the newly selected Master, but because PADT was sent to clients the recovery time is faster (no need to wait for PPPoE session timeout). On the network side (towards the LNS) an existing tunnel towards the LNS can be used to re-establish the sessions or in case that none exists, a new tunnel will be established. Then there is no need for a redundant interface.

Note: The L2TP tunnel carrying the sessions must always be terminated on the Master LAC.

In case of nodal failure, the sessions within the old tunnel on the LAC times out (CDN cannot be sent from the new Master since there is no tunnel state preserved across redundant LAC nodes). During the time-out period, the LNS must maintain double the amount of failed sessions (stale ones plus the new ones). This model is shown in Figure 44.

Figure 44 **Non-Hitless Interface/Node Protection on the LAC**

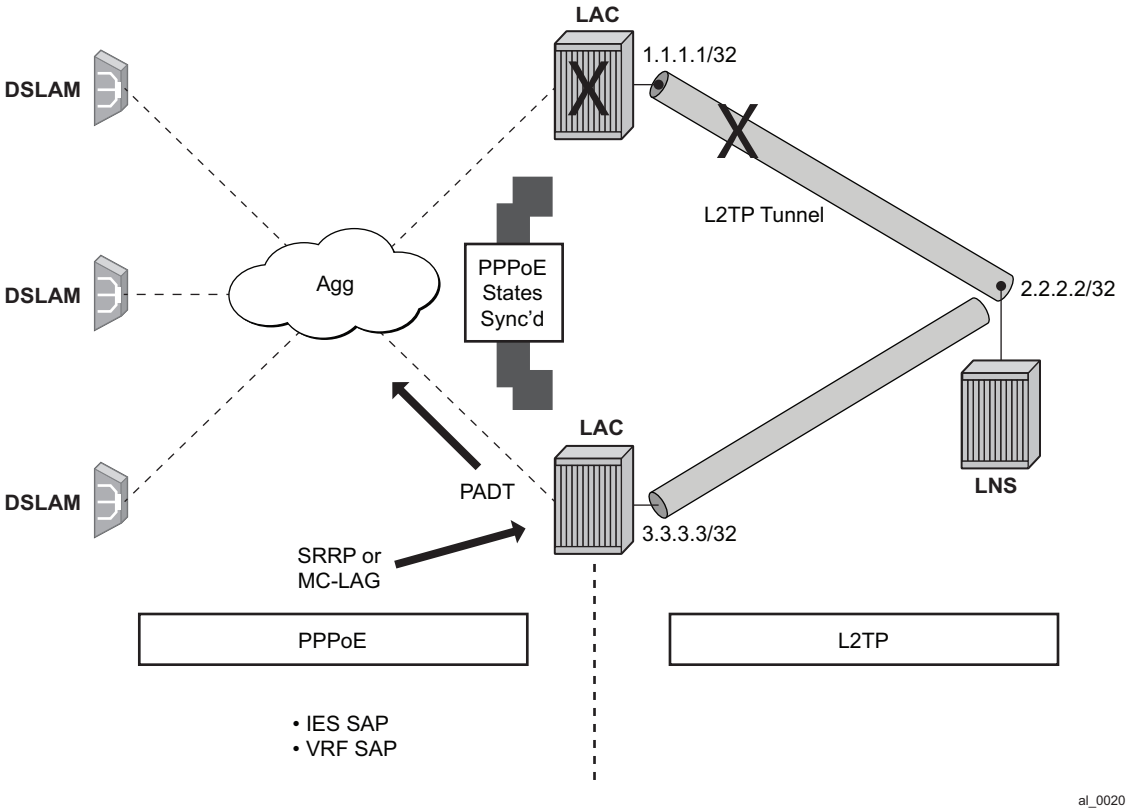

6.3.1 Per-ISP Egress L2TP DSCP Reclassification

Wholesale providers can deliver Internet access to directly connected PPP users through third party ISPs. This involves the users connecting to an L2TP Access Concentrator (LAC) with their traffic being tunneled to and from an L2TP Network Server (LNS) in their ISP.

If there is a requirement to support per-ISP (and per-subscriber host) QOS control for downstream traffic on the LAC towards the users based on the DSCP marking in the L2TP header, the command **use-ingress-l2tp-dscp** must be configured within the sla-profile selected for the users. An example topology is shown in Figure 45.

Figure 45 **ISP Internet Access through Wholesale Provider**

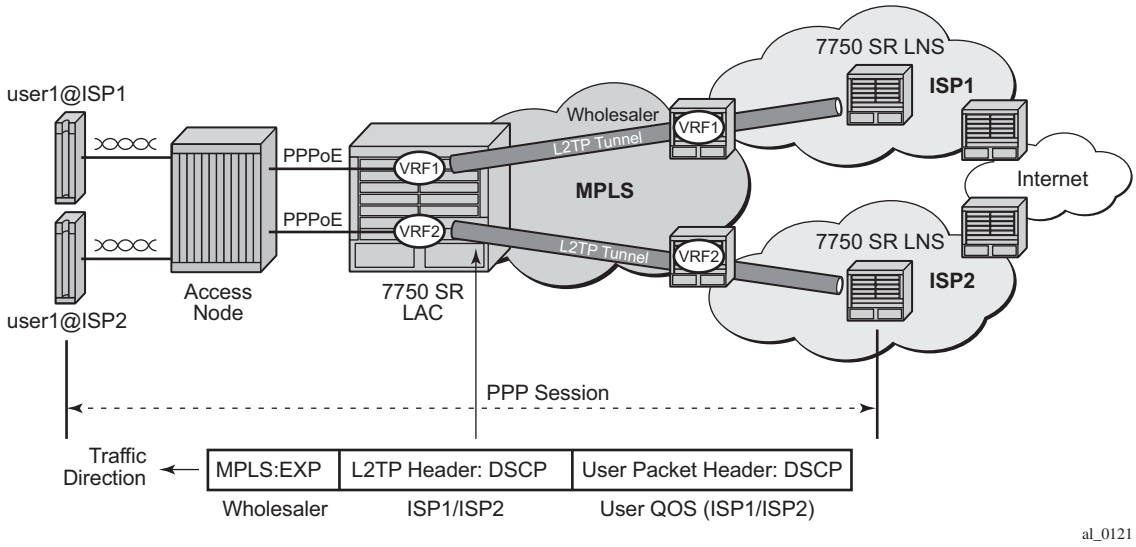

The downstream traffic arrives at the LAC with:

- An MPLS header (because of the VRF encapsulation). This contains EXP bits which are set based on the wholesale provider's QOS scheme.
- An L2TP header (because of the L2TP tunnel to the ISP). This contains DSCP bits in its IP header which are set by the originating ISP.
- A user IP packet header. This contains DSCP bits which could be set by the ISP or by the originating Internet application.

The network ingress on the LAC would normally use the MPLS EXP bits for traffic QoS classification, however, this matches the wholesale provider's QoS scheme.

It is possible to apply the **ler-use-dscp** parameter at the LAC network ingress to classify based on the L2TP header DSCP, but this would require the QoS schemes used by all ISPs, and the wholesale provider, to have a consistent interpretation of the DSCP bits.

If the standard egress IP reclassification is used, the QOS would be dependent on the DSCP in the user packet.

Configuring the parameter **use-ingress-l2tp-dscp** in the sla-profile of the ISP1 and ISP2 users will force the egress QoS control to be based on the DSCP from the L2TP header received on the LAC (which is set by ISP1/ISP2). This provides per-ISP (and per-subscriber host) QoS control for downstream traffic on the LAC towards the users.

6.4 Traffic Steering on L2TP LAC

Traffic steering on L2TP LAC allows wholesale providers to forward L2TP-encapsulated packets going to and coming from LNS to Value-Added Services (VAS).

Traffic steering on L2TP LAC consists of the following components:

- Steering profile — contains steering configuration (Access/VAS routers and next hop) information that is applied to the subscriber host for the PPPoE/L2TP session
- Steered traffic — all PPP packets to and from the PPPoE host that have a steering profile attached are forwarded through VAS.

  PPP packets include LCP/NCP control packets, LCP echo and echo reply, and user data packets.
- Non-steered traffic — consists of packets for the L2TP control channel and PPP packets of the subscriber host that don't have a steering profile
- Base Router (can also be a VPRN) — the routing instance terminates subscriber host and L2TP tunnels or sessions to LNS
- Access VAS Router — the routing instance forwards upstream (to LNS) steered traffic to VAS and receives downstream (to subscriber) steered traffic from VAS.

  An access VAS router needs to be a VPRN. A base routing instance cannot be specified as an access VAS router.
- Network VAS Router (optional) — the routing instance receives upstream (to LNS) steered traffic from VAS and forwards downstream (to subscriber) steered traffic to VAS.

  The creation of a network VAS router is not mandatory and a base router can also perform the same function.
- Network VAS next-hop — the next-hop IP address to reach VAS from the Network VAS Router or Base Router for downstream traffic. This address must be specified in the steering profile.

Figure 46 shows traffic steering on L2TP LAC.

Figure 46 **Traffic Steering on L2TP LAC**

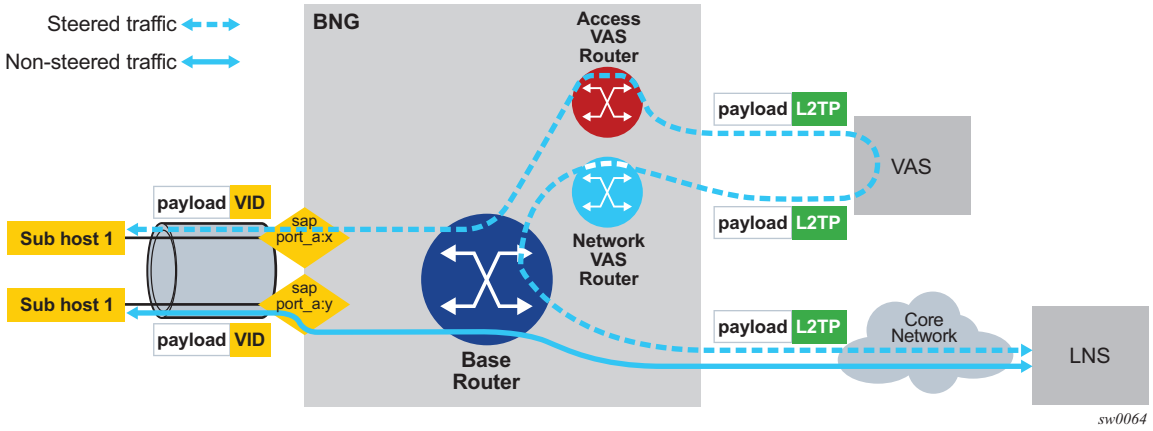

6.4.1 Steering Activation and Deactivation

Traffic steering can be activated by the LUDB and RADIUS Access-Accept/CoA messages that include the Alc-Steering-Profile VSA.

Steering can be deactivated by a RADIUS CoA that includes an Alc-Remove-Override VSA to remove Alc-Steering-Profile generated by an AAA server or BNG node itself using a CLI command.

The following CLI example shows the activation of a steering profile:

```
tools perform subscriber-mgmt coa alc-subscr-id subscriber-
1 attr evs,241,6527,25=”steering-profile-1”
```

The following CLI example shows the deactivation of a steering profile:

```
tools perform subscriber-mgmt coa alc-subscr-id subscriber-
1 attr 6527,238="deactivate 241.26.6527.25"
```

6.4.2 Steering States

Each PPPoE/L2TP session has an operational state of traffic steering that prevents a traffic black hole caused by problems such as an incorrectly configured network or a temporary VAS outage.

The steering states are:

- Non-steered — a steering profile is not applied to the session and traffic steering is not performed
- Steered — a steering profile is applied to the session and traffic steering is performed
- Steering-failure — a steering profile is applied to the session but traffic steering is not performed, possibly due to a misconfiguration of the steering profile, L2TP, or IP routing

As the conditions to perform traffic steering are met or lost, the steering state transitions in and out of steering failure.

6.4.3 Configuring Traffic Steering on L2TP LAC

The following steps show the commands used to configure traffic steering on L2TP LAC.

1. Enable L2TP on a base router.

```
router
    l2tp
         no shutdown
         exit
    exit
```

2. Create VAS routers and interfaces.

```
vprn 100 customer 1 create
    description "VAS router for access"
    route-distinguisher 65001:100
    vrf-target target:65001:100
    l2tp
        no shutdown
        exit
    interface "L2TP LAC endpoint"
        address 10.0.0.1
        loopback
        exit
    interface "Access interface to VAS"
        address 10.0.100.100/24
        vas-if-type to-from-access
        exit
    static-route-entry 20.0.0.1/32 next-hop 10.0.100.200
vprn 200 customer 1 create
    description "VAS router for network"
    route-distinguisher 65001:200
    vrf-target target:65001:200
    interface "Network interface to VAS"
        address 10.0.100.200/24
        exit
    grt-lookup
        enable-grt
```

```
        exit
    static-route-entry 10.0.0.1/32 next-hop 10.0.100.100
```

The L2TP endpoint IP addresses used by the base router and the access router must be the same.

The VAS-facing interface on the access router must be configured as **vas-if-type to-from-access** to avoid a traffic loop between BNG and VAS.

3. Configure the steering profile.

```
subscriber-mgmt
    steering-profile "SP1" create
        access router 100
            exit
        network next-hop 10.0.100.200 router 200
            exit
    exit
```

4. (Optional) Configure the LUDB host entry for the steered L2TP/PPPoE session.

```
subscriber-mgmt
    local-user-db "LAC-steering"
        ppp
            match-list sap-id
            host "lag-1:1" create
                host-identification
                    sap-id "lag-1:1"
                exit
                auth-policy "AUTH1"
                pre-auth-policy "PRE1"
                steering-profile "SP1"
                identification-strings 254 create
                    sla-profile-string "SLA1"
                    sub-profile-string "SUB1"
                exit
                no shutdown
            exit
```

6.5 L2TP Tunnel RADIUS Accounting

Figure 47 shows an L2TP tunnel RADIUS accounting configuration.

Figure 47 **L2TP Tunnel Accounting**

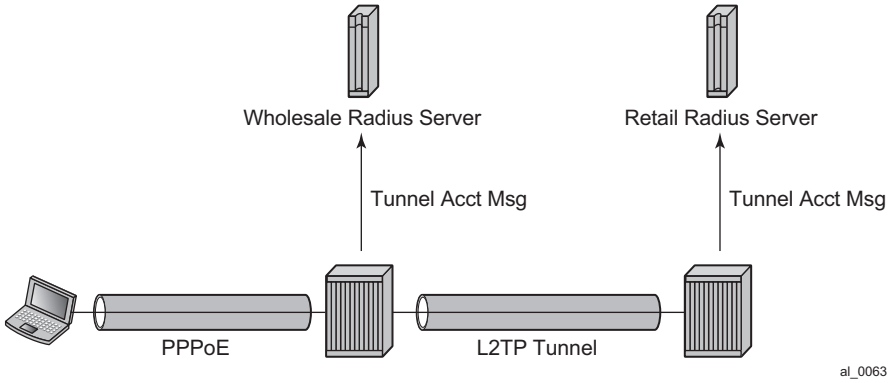

When L2TP tunnel accounting is enabled, except for **host** or **sla-profile**-based accounting packets and attributes, the following are additional accounting packets and attributes:

- Accounting packets:
  - tunnel-start/stop/reject
  - tunnel-link-start/stop/reject

  There are no interim updates for L2TP tunnel/session accounting.
- RADIUS accounting attributes:
  - Tunnel-Assignment-Id (LAC only)
  - Acct-Tunnel-Connection
  - Acct-Tunnel-Packets-Lost

These attributes were added into current account-start/stop/interim-update packets (host accounting/sla-profile accounting)

Tunnel level accounting and session level accounting can be enabled or disabled independently.

New accounting packets and related RADIUS attribute list are described in Table 68.

Some considerations of RADIUS attributes are described in RADIUS Attributes Value Considerations.

6.5.1 Accounting Packets List

Table 68 describes L2TP tunnel accounting behavior along with some key RADIUS attributes (apply for both LAC and LNS):

Table 68 **L2TP Tunnel Accounting Behavior**

Act-Packet	When	Key Attributes
Tunnel-Start	A new L2TP tunnel is created	Acct-Session-ID Event-Timestamp Tunnel-Type:0 Tunnel-Medium-Type:0 Tunnel-Assignment-Id:0 Tunnel-Client-Endpoint:0 Tunnel-Client-Auth-Id:0 Tunnel-Server-Endpoint:0 Tunnel-Server-Auth-Id:0
Tunnel-Reject	A new L2TP tunnel creation failed	Acct-Session-Id Event-Timestamp Tunnel-Type:0 Tunnel-Medium-Type:0 Tunnel-Assignment-Id:0 Tunnel-Client-Endpoint:0 Tunnel-Client-Auth-Id:0 Tunnel-Server-Endpoint:0 Acct-Terminate-Cause

Table 68 **L2TP Tunnel Accounting Behavior (Continued)**

Act-Packet	When	Key Attributes
Tunnel-Stop	An established L2TP tunnel is removed	Acct-Session-Id Event-Timestamp Tunnel-Type:0 Tunnel-Type:0 Tunnel-Medium-Type:0 Tunnel-Assignment-Id:0 Tunnel-Client-Endpoint:0 Tunnel-Client-Auth-Id:0 Tunnel-Server-Endpoint:0 Tunnel-Server-Auth-Id:0 Tunnel-Server-Auth-Id:0 Acct-Session-Time Acct-Input-Gigawords Acct-Input-Octets Acct-Output-Gigawords Acct-Output-Octets Acct-Output-Octets Acct-Input-Packets Acct-Output-Packets Acct-Terminate-Cause
Tunnel-Link-Start	An L2TP session is created	User-Name Acct-Session-Id — This is the same as Acct-Session-id in access-request of host auth Event-Timestamp Service-Type — Framed Class Tunnel-Type:0 Tunnel-Medium-Type:0 Tunnel-Assignment-Id:0 Tunnel-Client-Endpoint:0 Tunnel-Client-Auth-Id:0 Tunnel-Server-Endpoint:0 Tunnel-Server-Auth-Id:0 Acct-Tunnel-Connection — See RADIUS Attributes Value Considerations

Table 68 **L2TP Tunnel Accounting Behavior (Continued)**

Act-Packet	When	Key Attributes
Tunnel-Link-Reject	A new L2TP session creation is failed	Acct-Session-Id — Should be as same as Acct-Session-id in access-request of host auth Event-Timestamp Tunnel-Type:0 Tunnel-Medium-Type:0 Tunnel-Assignment-Id:0 Tunnel-Client-Endpoint:0 Tunnel-Client-Auth-Id:0 Tunnel-Server-Endpoint:0 Acct-Terminate-Cause Acct-Tunnel-Connection
Tunnel-Link-Stop	A established L2TP session is removed	User-Name Acct-Session-Id — Should be as same as Acct-Session-id in access-request of host auth Event-Timestamp Service-Type — Framed Class Tunnel-Type:0 Tunnel-Medium-Type:0 Tunnel-Assignment-Id:0 Tunnel-Client-Endpoint:0 Tunnel-Client-Auth-Id:0 Tunnel-Server-Endpoint:0 Tunnel-Server-Auth-Id:0 Acct-Tunnel-Connection Acct-Session-Time Acct-Input-Gigawords Acct-Input-Octets Acct-Output-Gigawords Acct-Output-Octets Acct-Input-Packets Acct-Output-Packets Acct-Tunnel-Packets-Lost Acct-Terminate-Cause

Notes:

- Errors will occur if there are multiple hosts sharing the same sla-profile instance and then these hosts go to different tunnel.
- 7750 SRs have an internal limitation of 500 pps for accounting messages. This feature shares the same limitation

6.5.2 RADIUS Attributes Value Considerations

- The value of Acct-Tunnel-Connection uniquely identifies an L2TP session. In order to match LAC and LNS accounting records, the value of Acct-Tunnel-Connection is determined by a method shared by LAC and LNS. This means for a given L2TP session, Acct-Tunnel-Connection from the LAC and LNS are the same.
- Current ESM statistics are used in tunnel link and tunnel level accounting. This applies to the standard attribute and the 7750 SR's own VSA.
- Tunnel level accounting statistics must aggregate all session statistics that belong to the tunnel.

Note: There could be sessions come and go before tunnel is down, so system need to remember the stats of every session that has been created within the tunnel.

This applies for both standard attribute and 7750 SR's own VSA.

- The value of Acct-Tunnel-Packets-Lost is the aggregation of all discarded packets on both ingress and egress.
- L2TP tunnel accounting on LAC can be enabled from RADIUS using the Alc-Tunnel-Acct-Policy VSA. This attribute overrides the locally configured L2TP RADIUS accounting policy.
- Refer to the *7750 SR and VSR RADIUS Attributes Reference Guide* for attributes that are applicable in RADIUS L2TP tunnel accounting.

6.6 MLPPP on the LNS Side

With MLPPP, the counter on LNS side is only available for the bundle, not for each link, so the SR OS's behavior is:

- For each new link session system sends a tunnel-link-start.
- For each link session that is deleted system sends a tunnel-link-stop.
- For all link sessions except the last one system reports 0 for all counters.
- For the last link session, system reports the actual counters for the bundle.

6.7 LNS Reassembly

LNS reassembly is supported in the BB-ISA. Fragments are collected and reassembled. Once the entire L2TP packet is reassembled, the packet will either be decapsulated or sent to the CPM without change.

The delivery of the L2TP packets to the BB-ISA depends on the certain fields in the L2TP header. The forwarding decision on the ingress LNS side in the upstream direction (LAC->LNS) is based on the tunnel-id/session-id combination and the T-bit (message type bit – control or data) in L2TP header.

Control type messages are delivered directly to the CPM. CPM performs L2TP decapsulation and processes the message (tunnel or session setup/teardown related messages or tunnel hellos). The CPM provides forwarding information to the forwarding plane (ingress/egress IOM and the carrier IOM) and to the BB-ISA (tunnel-src plus tunnel-id/session-id plus generated-mac-addr and SAP).

Data type messages are delivered directly to the BB-ISA. The BB-ISA decapsulates the L2TP packets and forwards them to the carrier IOM as a quasi-PPPoE frame (ESM forwarding module).

Since the LAC fragments the packets in the upstream direction, the L2TP header is preserved only in the first fragment. Therefore, the crucial forwarding information needed by LNS is lost in all consecutive fragments. If a fragments ends up in the wrong BB-ISA with no reassembly context for the fragment, the fragment will be dropped.

Similarly, the information whether to forward the fragment to the BB-ISA (data packet) or the CPM (control packet) is lost.

In order to support LSN reassemble, the following configuration limitations are imposed:

- Only one pair of active/standby BB-ISAs are supported. This way all fragments will be forwarded to the same active BB-ISA that maintains all reassembly contexts for all fragments.
- All fragments, regardless of the packet type, are forwarded to the active BB-ISA. Once the L2TP packet is reassembled, it will be determined whether the packet is:
  - A data packet — The packet will be decapsulated and a quasi PPPoE packet will be forwarded to the carrier IOM (ESM function).
  - A control packet — The packet will not be decapsulated but instead it will be forwarded as L2TP packet to the CPM.

The **lns-reassembly** commands that inform the ingress forwarding plane that all L2TP packets should be sent to the BB-ISA are configured in the **config>router>l2tp** and **config>service>vprn>l2tp** contexts.

6.8 LNS Subscriber Policers

6.8.1 Policer Support

LNS subscribers' SLA profiles support policers. In the egress direction, HQoS manageable policers are also supported.

The following QoS features are supported for LNS subscribers:

- ingress policer with h-pol
- egress policer with h-pol
- policer-hqos-manageable
- policer-output-queue
- egress queuing, policer to local queue
- egress queuing, bypass policers
- egress queuing, flow based

6.9 L2TP Command Reference

6.9.1 Command Hierarchies

- L2TP Configuration Commands
- Steering Profile Commands
- L2TP Tunnel RADIUS Accounting Commands
- Show Commands
- Clear Commands
- Debug Commands
- Tools Commands

6.9.1.1 L2TP Configuration Commands

```
configure
    — router
        — l2tp
            — df-bit-lac {always | never}
            — no df-bit-lac
            — group tunnel-group-name [create]
            — no group tunnel-group-name
                — df-bit-lac {always | never | default}
                — no df-bit-lac
                — ppp
                    — default-group-interface ip-int-name service-id service-id
                    — default-group-interface ip-int-name service-name svc-name
                    — no default-group-interface
                    — [no] lcp-force-ack-accm
                — tunnel tunnel-name [create]
                — no tunnel tunnel-name
                    — df-bit-lac {always | never | default}
                    — no df-bit-lac
                    — ppp
                        — default-group-interface ip-int-name service-id service-
                          id
                        — default-group-interface ip-int-name service-name svc-
                          name
                        — no default-group-interface
                        — lcp-force-ack-accm {always | never}
                        — no lcp-force-ack-accm
            — ignore-avps [sequencing-required]
            — no ignore-avps
            — tunnel-selection-blacklist
```

```
            — add-tunnel never
            — add-tunnel on reason [reason]
            — no add-tunnel
            — max-list-length unlimited
            — max-list-length count
            — no max-list-length
            — max-time minutes
            — no max-time
            — no timeout-action
            — timeout-action action

configure
    — system
        — l2tp
            — non-multi-chassis-tunnel-id-range start l2tp-tunnel-id end l2tp-tunnel-id
            — non-multi-chassis-tunnel-id-range default
            — no non-multi-chassis-tunnel-id-range

configure
    — redundancy
        — multi-chassis
            — peer
                — sync
                    — track-srrp-instances
                        — [no] track-srrp [srrp-instance]
                            — l2tp-tunnel-id-range start l2tp-tunnel-id end
                                l2tp-tunnel-id
                            — no l2tp-tunnel-id-range

configure
    — router
        — l2tp
            — failover
                — recovery-max-session-lifetime minutes
                — no recovery-max-session-lifetime
                — recovery-method method
                — no recovery-method
                — recovery-time seconds
                — no recovery-time
                — track-srrp srrp-instance peer ip-address sync-tag tag
                — [no] track-srrp srrp-instance
            — group tunnel-group-name [create]
            — no group tunnel-group-name
                — failover
                    — recovery-method method
                    — no recovery-method
                    — recovery-time seconds
                    — no recovery-time
                — tunnel tunnel-name [create]
                — no tunnel tunnel-name
                    — failover
                        — recovery-method method
                        — no recovery-method
```

```
                                    — recovery-time seconds
                                    — no recovery-time

configure
    — service
        — vprn
            — l2tp
                — group tunnel-group-name [create]
                — no group tunnel-group-name
                    — ppp
                        — default-group-interface ip-int-name service-id service-id
                        — default-group-interface ip-int-name service-name svc-name
                        — no default-group-interface
                        — [no] lcp-force-ack-accm
                    — tunnel tunnel-name [create]
                    — no tunnel tunnel-name
                        — ppp
                            — default-group-interface ip-int-name service-id service-id
                            — default-group-interface ip-int-name service-name svc-name
                            — no default-group-interface
                            — lcp-force-ack-accm {always | never}
                            — no lcp-force-ack-accm
                — ignore-avps [sequencing-required]
                — no ignore-avps
                — failover
                    — recovery-max-session-lifetime minutes
                    — no recovery-max-session-lifetime
                    — recovery-method method
                    — no recovery-method
                    — recovery-time seconds
                    — no recovery-time
                    — track-srrp srrp-instance peer ip-address sync-tag tag
                    — [no] track-srrp srrp-instance
```

6.9.1.2 Steering Profile Commands

```
configure
    — subscriber-mgmt
        — steering-profile steering-profile-name [create]
        — no steering-profile steering-profile-name
            — access router router-instance
            — no access
            — description description-string
            — no description
            — network next-hop ip-address [router router-instance]
            — no network
```

6.9.1.3 L2TP Tunnel RADIUS Accounting Commands

```
configure
    — aaa
        — l2tp-accounting-policy policy-name [create]
        — no l2tp-accounting-policy policy-name
            — accounting-type [session] [tunnel]
            — no accounting-type
            — description description-string
            — no description
            — include-radius-attribute
                — [no] nas-identifier
                — nas-port binary-spec
                — no nas-port
                — nas-port-id
                — nas-port-id [prefix-string string] [suffix suffix-option]
                — no nas-port-id
                — nas-port-type [type]]
                — no nas-port-type
            — radius-accounting-server
                — access-algorithm {direct | round-robin}
                — no access-algorithm
                — retry count
                — no retry
                — router service-name service-name
                — router router-instance
                — no router
                — server server-index address ip-address secret key [hash | hash2]
                    [port port]
                — no server server-index
                — source-address ip-address
                — no source-address
                — timeout [sec seconds] [min minutes]
                — no timeout
            — request-script-policy script-policy
            — no request-script-policy
```

6.9.1.4 Show Commands

```
show
    — subscriber-mgmt
        — ppp-policy [ppp-policy-name [association]]
    — service
        — id service-id
            — pppoe
```

```
— session [interface ip-int-name | ip-address | sap sap-id] [type pppoe-session-type] [session-id session-id] [mac ieee-address] [ip-address ip-address[/mask]] [port port-id] [no-inter-dest-id | inter-dest-id intermediate-destination-id] [steering-profile steering-profile] [router-advertisement-policy policy-name] [detail | statistics]
— session l2tp-connection-id connection-id [detail | statistics]
— statistics [{sap sap-id | interface ip-int-name | ip-address}
— summary
— system
— l2tp
— redundancy
— multi-chassis
— sync [peer ip-address] [statistics]
— sync peer ip-address detail
```

6.9.1.5 Clear Commands

```
clear
— service
— id service-id
— pppoe
— session all [no-padt]
— session {interface ip-int-name | ip-address | sap sap-id} [mac ieee-address] [session-id session-id] [ip-address ip-address[/mask]] [port port-id] [no-inter-dest-id | inter-dest-id intermediate-destination-id] [no-padt]
— statistics [{sap sap-id | interface ip-int-name | ip-address}]
```

6.9.1.6 Debug Commands

```
debug
— service
— id service-id
— [no] ppp
— [no] event
— dhcp-client [terminate-only]
— no dhcp-client
— ppp [terminate-only]
— no ppp
— [no] mac ieee-address
— [no] packet
— detail-level {low | medium | high}
— no detail-level
— [no] dhcp-client
— discovery [padi] [pado] [padr] [pads] [padt]
— no discovery
```

```
                        — mode {dropped-only | ingr-and-dropped |egr-ingr-and-
                                dropped}
                        — no mode
                        — ppp [lcp] [pap] [chap] [ipcp]
                        — no ppp
                    — [no] sap sap-id
— router [router-instance]
    — [no] l2tp
        — assignment-id assignment-id
            — [no] event
                — [no] recovery
                — [no] recovery-failed
        — [no] event
            — [no] recovery
            — [no] recovery-failed
        — group tunnel-group-name
            — [no] event
                — [no] recovery
                — [no] recovery-failed
        — peer ip-address [udp-port port]
            — [no] event
                — [no] recovery
                — [no] recovery-failed
        — tunnel connection-id
            — [no] event
                — [no] recovery
                — [no] recovery-failed
```

6.9.1.7 Tools Commands

```
tools
    — perform
        — subscriber-mgmt
            — edit-ppp-session sap sap-id [user-name user-name] [mac mac-address]
                [session-id session-id] [ip-address ip-address] [subscriber sub-ident-
                string] [sub-profile-string sub-profile-string] [sla-profile-string sla-profile-
                string] [inter-dest-id intermediate-destination-id] [ancp-string ancp-string]
                [app-profile-string app-profile-string]
            — edit-ppp-session svc-id service-id [user-name user-name] [mac mac-
                address] [session-id session-id] [ip-address ip-address] [subscriber sub-
                ident-string] [sub-profile-string sub-profile-string] [sla-profile-string sla-
                profile-string] [inter-dest-id intermediate-destination-id] [ancp-string ancp-
                string] [app-profile-string app-profile-string]
            — eval-lease-state [svc-id service-id] [sap sap-id] [subscriber sub-ident-string]
                [ip ip-address] [mac ieee-address]
            — local-user-db local-user-db-name [mac ieee-address]
                — ipoe
```

— **host-lookup** [**mac** *ieee-address*] [**remote-id** *remote-id-ascii*] [**sap-id** *sap-id*] [**service-id** *service-id*] [**string** *vso-string*] [**system-id** *system-id*] [**option60** *option-60-ascii*] [**circuit-id** *circuit-id-ascii*] [**circuit-id-hex** *circuit-id-hex*] [**option60-hex** *option60-hex*] [**remote-id-hex** *remote-id-hex*] [**derived-id** *derived-id*] [**ip-prefix** *ip-prefix/ip-prefix-length*]

— **ppp**

— **authentication password** *password* [**mac** *ieee-address*] [**remote-id** *remote-id*] [**circuit-id** *circuit-id*] **user-name** *user-name* [**service-name** *service-name*]

— **authentication password** *password* [**mac** *ieee-address*] [**remote-id** *remote-id*] [**circuit-id-hex** *circuit-id-hex*] **user-name** *user-name* [**service-name** *service-name*]

— **host-lookup** [**circuit-id** *circuit-id*] [**circuit-id-hex** *circuit-id-hex*] [**derived-id** *derived-id*] [**mac** *ieee-address*] [**remote-id** *remote-id*] [**remote-id-hex** *remote-id-hex*] [**sap-id** *sap-id*] [**service-name** *service-name*] [**user-name** *user-name*]

6.9.2 Command Descriptions

6.9.2.1 L2TP Configuration Commands

6.9.2.1.1 Global Commands

description

Syntax **description** *description-string*
no description

Context config>aaa>l2tp-acct-plcy

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Parameters *description-string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

shutdown

Syntax [**no**] **shutdown**

Context config>aaa>l2tp-acct-plcy

Description This command administratively disables an entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they may be deleted.

The **no** form of this command places the entity into an administratively enabled state.

6.9.2.1.2 L2TP Commands

next-attempt

Syntax **next-attempt** {**same-preference-level** | **next-preference-level**}
no next-attempt

Context config>router>l2tp
config>service>vprn>l2tp

Description This command enables tunnel selection algorithm based on the tunnel preference level.

The **no** form of this command reverts to the default.

Default next-preference-level

Parameters **same-preference-level** — If the **tunnel-spec** selection algorithm evaluates into a tunnel that is currently unavailable (for example, a tunnel in a blacklist) then the next elected tunnel, if available, is chosen within the same preference-level as the last attempted tunnel. Only when all tunnels within the same preference level are exhausted, the tunnel selection algorithm moves to the next preference level.

In case that a new session setup request is received while all tunnels on the same preference level are blacklisted, the L2TP session tries to be established on blacklisted tunnels before the tunnel selection moves to the next preference level.

next-preference-level — If the **tunnel-spec** selection algorithm evaluates into a tunnel that is currently unavailable (for example tunnel in a blacklist) then the selection algorithm tries to select the tunnel from the next preference level, even though the tunnels on the same preference level might be available for selection.

Default next-preference-level

replace-result-code

Syntax **replace-result-code** *code* [*code*)]
no replace-result-code

Context config>router>l2tp
config>service>vprn>l2tp

Description This command replaces CDN Result-Code 4, 5 and 6 on LNS with the Result Code 2. This is needed for interoperability with some implementation of LAC which only take action based on CDN Result-Code 2, while ignore CDN Result-Code 4, 5 and 6.

The **no** form of this command reverts to the default.

Parameters *code* — Specifies up to three L2TP Result codes that need to be replaced.

Values cdn-tmp-no-facilities — CDN Result-Code 4 on LNS is replaced with the result code 2 before it is sent to LAC.

cdn-prem-no-facilities — CDN Result-Code 5 on LNS is replaced with the result code 2 before it is sent to LAC.

cdn-inv-dest — CDN Result-Code 6 on LNS is replaced with the result code 2 before it is sent to LAC.

df-bit-lac

Syntax **df-bit-lac** {**always** | **never**}
no df-bit-lac

Context config>router>l2tp
config>service>vprn>l2tp

Description By default, the LAC df-bit-lac is always set and sends all L2TP packets with the DF bit set to 1. The DF bit is configurable to allow downstream routers to fragment the L2TP packets. The LAC itself will not fragment L2TP packets. L2TP packets that have a larger MTU size than what the LAC egress ports allows are dropped.

The **no** form of this command reverts to the default.

Default df-bit-lac always

Parameters **always** — Specifies that the LAC sends all L2TP packets with the DF bit set to 1.

never — Specifies that the LAC sends all L2TP packets with the DF bit set to 0.

df-bit-lac

Syntax **df-bit-lac** {**always** | **never** | **default**}
no df-bit-lac

Context config>router/service>vprn>l2tp>group
config>router/service>vprn>l2tp>group>tunnel

Description By default, the LAC df-bit-lac is set to default and sends all L2TP packets with the DF bit set to 1. The DF bit is configurable to allow downstream routers to fragment the L2TP packets. The LAC will not fragment L2TP packets. L2TP packets that have a larger MTU size than what the LAC egress ports allows are dropped. The configuration of the df-bit can be overridden at different levels: l2tp, tunnel, and group. The configuration at the tunnel level overrides the configuration on both the group and l2tp levels. The configuration at the group level overrides the configuration on l2tp.

The **no** form of this command reverts to the default.

Default df-bit-lac default

Parameters **always** — Specifies that the LAC sends all L2TP packets with the DF bit set to 1.

never — Specifies that the LAC sends all L2TP packets with the DF bit set to 0.

default — Follows the DF-bit configuration specified on upper levels.

ignore-avps

Syntax **ignore-avps** [**sequencing-required**]
no ignore-avps

Context config>router>l2tp
config>service>vprn>l2tp

Description This command specifies the L2TP AVPs that should be ignored in L2TP session control.

The **no** form of this command reverts to the default.

Parameters **sequencing-required** — Ignores the [39] Sequencing Required AVP on LNS when present in the L2TP ICCN message received from LAC. By default, the session at LNS would be disconnected, in this case with the Call Disconnect Notify (CDN) error code unknownMandatoryReceive(8). Note that when configured, to ignore the Sequencing Required AVP there is no Sequence Numbers inserted into the data channel.

group

Syntax **group** *tunnel-group-name* [**create**]
no group *tunnel-group-name*

Context config>router>l2tp
config>service>vprn>l2tp

Description This command configures an L2TP tunnel group.

The **no** form of this command reverts to the default.

Parameters *tunnel-group-name* — Specifies a name string to identify a L2TP group up to 63 characters.

create — This keyword is mandatory when creating a tunnel group name. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

tunnel

Syntax **tunnel** *tunnel-name* [**create**]
no tunnel *tunnel-name*

Context config>router>l2tp>group
config>service>vprn>l2tp>group

Description This command configures an L2TP tunnel. A tunnel exists between a LAC-LNS pair and consists of a Control Connection and zero or more L2TP sessions. The tunnel carries encapsulated PPP datagrams and control messages between the LAC and the L2TP Network Server (LNS).

The **no** form of this command removes the tunnel name from the configuration.

Parameters *tunnel-name* — Specifies a valid string to identify an L2TP, up to 32 characters.

create — Mandatory keyword to create a new tunnel.

default-group-interface

Syntax **default-group-interface** *ip-int-name* **service-id** *service-id*
default-group-interface *ip-int-name* **service-name** *svc-name*
no default-group-interface

Context config>router>l2tp>group>ppp
config>router>l2tp>group>tunnel>ppp
config>service>vprn>l2tp>group>ppp
config>service>vprn>l2tp>group>tunnel>ppp

Description This command configures the group interface where the PPP sessions are established when the authentication server does not specify the group interface.

The **no** form of this command removes the interface name or service ID from the configuration.

Parameters *ip-int-name* — Specifies an IP interface name up to 32 characters.

service-id *service-id* — Specifies an existing service identification number.

This variant of the command is only supported in 'classic' configuration-mode (**configure system management-interface configuration-mode classic**). The **default-group-interface** *ip-int-name* **service-name** *svc-name* variant can be used in all configuration modes.

Values {*id* | *svc-name*}

id:	1 to 2147483647
svc-name:	up to 64 characters (*svc-name* is an alias for input only. The *svc-name* gets replaced with an id automatically by SR OS in the configuration).

service-name *svc-name* — Specifies an existing service name up to 64 characters.

lcp-force-ack-accm

Syntax [**no**] **lcp-force-ack-accm**

Context config>router>l2tp>group>ppp
config>service>vprn>l2tp>group>ppp

Description This command enables or disables the LCP Asynchronous Control Character Map (ACCM) configuration option. When enabled, the LCP ACCM configuration option is acknowledged during LCP negotiation between the LNS and the PPP client. The option is then ignored and no ACCM mapping is done.

By default, the LCP ACCM configuration option is rejected.

The **no** form of this command reverts to the default.

ppp

Syntax **ppp**

Context config>router>l2tp>group

Description This command configures PPP for the L2TP tunnel group.

ppp

Syntax **ppp**

Context config>router>l2tp>group>tunnel

Description This command configures PPP for the L2TP tunnel

lcp-force-ack-accm

Syntax **lcp-force-ack-accm** {**always** | **never**}
no lcp-force-ack-accm

Context config>router>l2tp>group>tunnel>ppp
config>service>vprn>l2tp>group>tunnel>ppp

Description This command enables or disables the LCP Asynchronous Control Character Map (ACCM) configuration option. When enabled, the LCP ACCM configuration option is acknowledged during LCP negotiation between the LNS and the PPP client. The option is then ignored and no ACCM mapping is done.

By default, an L2TP tunnel inherits the configuration from the L2TP group CLI context.

The **no** form of this command reverts to the default.

Parameters **always** — Specifies to acknowledge the LCP ACCM configuration option, but not to perform ACCM mapping. This command overrides the group level configuration.

never — Specifies to reject the LCP ACCM configuration option. This command overrides the group level configuration.

tunnel-selection-blacklist

Syntax **tunnel-selection-blacklist**

Context config>router>l2tp

Description This command enables the context to configure L2TP Tunnel Selection Blacklist parameters.

add-tunnel

Syntax **add-tunnel never**
add-tunnel on *reason* [*reason*]
no add-tunnel

Context config>router>l2tp>tunnel-selection-blacklist
config>service>vprn>l2tp>tunnel-selection-blacklist

Description This command forces the tunnel to the blacklist and renders it unavailable for new sessions for the duration of pre-configured time. Peers are always forced to the blacklist in case that they time out (failure to receive response to control packets). In addition to time outs, certain events can be used to trigger placement of the tunnel on the blacklist.

The **no** form of this command reverts to the default.

Default add-tunnel never

Parameters *reason* — Specifies up to eight return codes or events that determine which tunnels are added to the blacklist.

Values

cdn-err-code — A tunnel is forced to the blacklist in case that CDN message with the Result Code 2 (Call disconnected for the reasons indicated in error code) is received.

cdn-inv-dest — A tunnel is forced to the blacklist in case that CDN message with the Result Codes 6 (Invalid destination) is received.

cdn-tmp-no-facilities — A tunnel is forced to the blacklist in case that CDN message with the Result Code 4 is received (Call failed due to lack of appropriate facilities being available - temporary condition) is received.

cdn-perm-no-facilities — A tunnel is forced to the blacklist in case that CDN message with the Result Codes 5 (Call failed due to lack of appropriate facilities being available - permanent condition) is received.

tx-cdn-not-established-in-time — A tunnel is forced to the blacklist in case that CDN message with the Result Code 10 (Call was not established within time allotted by LAC) is sent from the LAC to the LNS.

stop-ccn-err-code — A tunnel is forced to the blacklist in case that StopCCN message with the Result Code 2 (General error – Error Code indicates the problem) is sent or received.

stop-ccn-other — A tunnel is forced to the blacklist in case that StopCCN message with the following Result Codes is received:

(1) General request to clear control connection

(4) Requestor is not authorized to establish a control channel

(5) Protocol version not supported

(6) Requestor is being shutdown

or, in the case that the StopCCN with the following result codes is transmitted:

(4) Requestor is not authorized to establish a control channel.

(5) Protocol version not supported

The receipt of the following Result Codes never blacklists a tunnel:

(0) Reserved

(3) Control channel already exist

(7) Finite state machine error

(8) Undefined

Transmission of the following Result Codes never blacklists a tunnel:

(1) General request to clear control connection

(3) Control channel already exist

(6) Requestor is being shutdown

(7) Finite state machine error

addr-change-timeout — A timed-out tunnel for which the peer IP address has changed mid-session (from the one that is provided initially during configuration) is forced to the blacklist. In absence of this configuration option, only the configured peer for the tunnel is blacklisted, but not the tunnel itself which now has a different peer address than the one initially configured.

never — When specified, no tunnels are placed on blacklist under any circumstance. This parameter is available to preserve backward compatibility.

max-list-length

Syntax **max-list-length unlimited**
max-list-length *count*
no max-list-length

Context config>router>l2tp>tunnel-selection-blacklist
config>service>vprn>l2tp>tunnel-selection-blacklist

Description This command configured the maximum length of the peer/tunnel blacklist.

This command specifies how many items (tunnels or peers) can be in the tunnel-selection-blacklist. If a tunnel or peer needs to be added to the tunnel-selection-blacklist and the tunnel-selection-blacklist is full, the system will remove the item (tunnel or peer) from the blacklist that was in this blacklist for the longest time.

The **no** form of this command reverts to the default.

Default max-list-length unlimited

Parameters **unlimited** — Specifies there is no limit.

count — Specifies how many items (tunnels or peers) can be in the tunnel-selection-blacklist.

Values 1 to 65635

max-time

Syntax **max-time** *minutes*
no max-time

Context config>router>l2tp>tunnel-selection-blacklist

config>service>vprn>l2tp>tunnel-selection-blacklist

Description This command configures time for which an entity (peer or a tunnel) are kept in the blacklist.

The **no** form of this command reverts to the default.

Default max-time 5

Parameters *minutes* — Specifies the maximum time a tunnel or peer may remain in the blacklist.

Values 1 to 60

timeout-action

Syntax **timeout-action** *action*
no timeout-action

Context config>router>l2tp>tunnel-selection-blacklist
config>service>vprn>l2tp>tunnel-selection-blacklist

Description This command defines an action that is executed on the entity (peer/tunnel) in the blacklist once the entity becomes eligible for selection again.

The **no** form of this command reverts to the default.

Default timeout-action remove-from-blacklist

Parameters *action* — Specifies the Action to be taken when a tunnel or peer has been in the blacklist for the max-period of time.

Values **remove-from-blacklist** — The peer or tunnel in the blacklist is removed completely from the blacklist and made eligible for the selection process once the max-time expires. In this mode of operation, multiple new sessions can be mapped into the same, newly released tunnel from the blacklist. The first such session will try to setup the tunnel, while the other is buffered until the tunnel establishment process is completed. In case that the tunnel remains unavailable, it is placed in the blacklist again. Consequently all new sessions will have be re-negotiated over an alternate tunnel.

try-one-session — Once the max-time expired, the peer or tunnel in the blacklist is made available for selection only to a single new session request. Only upon successful tunnel establishment will the incoming new sessions be eligible to be mapped into this tunnel. This behavior will avoid session establishment delays in case that the tunnel just removed from the blacklist is still unavailable.

non-multi-chassis-tunnel-id-range

Syntax **non-multi-chassis-tunnel-id-range start** *l2tp-tunnel-id* **end** *l2tp-tunnel-id*
non-multi-chassis-tunnel-id-range default
no non-multi-chassis-tunnel-id-range

Context config>system>l2tp

Description This command sets the tunnel-id range that is used to allocate a new tunnel-id for a tunnel for which no multi-chassis redundancy is configured.

The **no** form of this command is a double negation and means all tunnel-IDs are configured for multi-chassis redundancy

Default Sets the tunnel-id range to the full tunnel-id range available on this system meaning that by default no tunnel-ID has multi-chassis redundancy.

non-multi-chassis-tunnel-id-range default or non-multi-chassis-tunnel-id-range start 1 end <maximum *tunnel-id*>

The default for **start** *l2tp-tunnel-id* is 1. No tunnel-ids are available for which no multi-chassis redundancy is configured when set to 0.

The default for **end** *l2tp-tunnel-id* is the maximum tunnel-id allowed on this system. The **end** *l2tp-tunnel-id* must be set to 0 when the **start** *l2tp-tunnel-id* is set to 0 and vice versa.

Parameters **start** *l2tp-tunnel-id* — Specifies the start of the range of L2TP tunnel identifiers that can be allocated by L2TP on this system, to be synchronized with Multi Chassis Redundancy Synchronization (MCS).

Values 0 to 16383

end *l2tp-tunnel-id* — Specifies the end of the range of L2TP tunnel identifiers that can be allocated by L2TP on this system, to be synchronized with Multi Chassis Redundancy Synchronization (MCS).

Values 1 to 16383

track-srrp-instances

Syntax **track-srrp-instances**

Context config>redundancy>multi-chassis>peer>sync

Description This command enables the context to configure tracked SRRP instances.

track-srrp

Syntax [**no**] **track-srrp** [*srrp-instance*]

Context config>redundancy>multi-chassis>peer>sync>track-srrp-instances

Description This command configures a tracked SRRP instance.

The **no** form of this command removes the SRRP instance identifier from the configuration.

Parameters *srrp-instance* — Indicates the unique identifier of the tracked SRRP instance.

Values 1 to 4294967295

l2tp-tunnel-id-range

Syntax **l2tp-tunnel-id-range start** *l2tp-tunnel-id* **end** *l2tp-tunnel-id*
no l2tp-tunnel-id-range

Context config>redundancy>multi-chassis>peer>sync>track-srrp-instances>track-srrp

Description This command sets the tunnel-id range that is used to allocate a new tunnel-id for a tunnel for which multi-chassis redundancy is configured to this MCS peer.

The **no** form of this command reverts to the default.

Parameters **start** *l2tp-tunnel-id* — Specifies the start of the range of L2TP tunnel identifiers that can be allocated by L2TP on this system, to be synchronized with Multi Chassis Redundancy Synchronization (MCS).

Values 1 to 16383

end *l2tp-tunnel-id* — Specifies the end of the range of L2TP tunnel identifiers that can be allocated by L2TP on this system, to be synchronized with Multi Chassis Redundancy Synchronization (MCS).

Values 1 to 16383

recovery-max-session-lifetime

Syntax **recovery-max-session-lifetime** *minutes*
no recovery-max-session-lifetime

Context config>router>l2tp>failover
config>service>vprn>l2tp>failover

Description This command configures the sub-set of sessions that this system attempts to synchronize in the Session State Synchronization phase as described in RFC 4951, Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP).

The **no** form of this command reverts to the default.

Default recovery-max-session-lifetime 2

Parameters *minutes* — Specifies the sub-set of sessions to recover.

Values 2 to 4294967295

recovery-method

Syntax **recovery-method** *method*
no recovery-method

Context config>router>l2tp>failover
config>service>vprn>l2tp>failover
config>router>l2tp>group>failover
config>service>vprn>l2tp>group>failover
config>router>l2tp>group>tunnel>failover
config>service>vprn>l2tp>group>tunnel>failover

Description This command sets the recovery method to be used for newly created tunnels.

The **no** form of this command reverts to the default.

Default recovery-method mcs on config>router>l2tp>failover and config>service>vprn>l2tp>failover

recovery-method default on config>router>l2tp>group>failover

recovery-method default on config>router>l2tp>group>tunnel>failover

recovery-method default on config>service>vprn>l2tp>group>failover

recovery-method default on config>service>vprn>l2tp>group>tunnel>failover

Parameters *method* — Describes how a pair of redundant LAC peers recover tunnel and session state (sequence numbers, for example) immediately after a failover.

Note: While failover is enabled, the tunnels and sessions proper are always kept synchronized between the redundant pair, regardless of the recovery method for the sequence numbers when a failover really occurs.

Values **mcs** — Specifies that the stateful information is recovered from the failover peer directly, using Multi-Chassis Redundancy Synchronization (MCS).

recovery-tunnel — Specifies that the stateful information is recovered as described in RFC 4951, *Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP)*. This method uses a recovery tunnel to the L2TP peer to pass the stateful information.

default — Specifies that the actual value must be derived from another object of the same type with a wider scope. Takes the value of the next higher level (not available in **config>router>l2tp>failover** and **config>service>vprn>l2tp>failover**).

recovery-time

Syntax **recovery-time** *seconds*
no recovery-time

Context config>router>l2tp>failover
config>service>vprn>l2tp>failover
config>router>l2tp>group>failover
config>service>vprn>l2tp>group>failover
config>router>l2tp>group>tunnel>failover
config>service>vprn>l2tp>group>tunnel>failover

Description This command sets the recovery time to be negotiated via RFC 4951. It represents the extra time this L2TP peer (LAC or LNS) needs to recover all its tunnels.

The **no** form of this command reverts to the default.

Default recovery-time 0 on config>router>l2tp>failover and config>service>vprn>l2tp>failover

Parameters *seconds* — The period of time, expressed in seconds, an endpoint asks its peer to wait before assuming the recovery process has failed.

Values 0 to 900

track-srrp

Syntax **track-srrp** *srrp-instance* **peer** *ip-address* **sync-tag** *sync-tag*
no track-srrp *srrp-instance*

Context config>router>l2tp>failover
config>service>vprn>l2tp>failover

Description This command sets the sync-tag to be used to synchronize the tunnels with track-srrp *srrp-id* to MCS *peer IP-@*. The same sync-tag should be configured on the MCS peer.

The **no** form of this command reverts to the default.

Default Removes the sync-tag for the indicated track-srrp.

Parameters *srrp-instance* — Specifies the Simple Router Redundancy Protocol (SRRP) instance used for Multi-Chassis redundancy failover that is associated with this Layer Two Tunneling Protocol Tunnel.

sync-tag — Specifies a synchronization tag to be used while synchronizing with the peer.

tunnel

Syntax **tunnel** *tunnel-name* [**create**]
no tunnel *tunnel-name*

Context config>router>l2tp>group

Description This command configures an L2TP tunnel.

The **no** form of this command removes the tunnel name from the group configuration.

Parameters *tunnel-name* — Specifies a string to identify a L2TP tunnel up to 32 characters.

create — Mandatory keyword to create a policy name.

6.9.2.2 Steering Profile Commands

steering-profile

Syntax **steering-profile** *steering-profile-name* [**create**]
no steering-profile *steering-profile-name*

Context config>subscr-mgmt

Description This command configures a steering profile mapping. A steering profile can be applied to each L2TP LAC subscriber host that requires traffic steering.

The **no** form of this command removes the specified steering profile.

Parameters *steering-profile-name* — Specifies the name of the steering profile.

access

Syntax **access router** *router-instance*
no access

Context config>subscr-mgmt>steering-profile

Description This command specifies a routing instance to be used as a network VAS router in the steering profile.

The **no** form of this command removes the router instance.

Parameters *router-instance* — Specifies the router instance to be used as an access VAS router.

Values
router-instance: *router-name* | *vprn-svc-id*
router-name: "Base"
vprn-svc-id: 1 to 2147483647

description

Syntax **description** *description-string*
no description

Context config>subscr-mgmt>steering-profile

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Parameters *description-string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

network

Syntax **network next-hop** *ip-address* [**router** *router-instance*]
no network

Context config>subscr-mgmt>steering-profile

Description This command specifies the downstream next-hop IP address and an optional routing instance to be used as a network VAS router in the steering profile.

The **no** form of this command removes the specified next-hop IP address and the router instance if specified.

Parameters *ip-address* — Specifies the IP address to be used as the downstream next-hop IP address in dotted decimal notation.

router-instance — Specifies the router instance to be used as an access VAS router.

Values
router-instance: *router-name* | *vprn-svc-id*
router-name: "Base"

vprn-svc-id: 1 to 2147483647

6.9.2.3 L2TP Tunnel RADIUS Accounting Commands

l2tp-accounting-policy

Syntax **l2tp-policy** *policy-name* [**create**]
no l2tp-policy

Context config>aaa

Description This command enables the L2TP accounting.

The **no** form of this command disables accounting.

Parameters *name* — The name of L2TP tunnel accounting policy.

create — Mandatory keyword to create a policy name.

accounting-type

Syntax **accounting-type** [**session**] [**tunnel**]
no accounting-type

Context config>aaa>l2tp-acct-plcy

Description This command specifies the accounting type for the L2TP tunnel accounting policy.

The **no** form of this command reverts to the default.

Default accounting-type session tunnel

Parameters **session** — Enables tunnel level accounting, including:
Tunnel-Link-Start
Tunnel-Link-Stop
Tunnel-Link-Reject

tunnel — Enables link level accounting, including:
Tunnel-Start
Tunnel-Stop
Tunnel-Reject

include-radius-attribute

Syntax [**no**] **include-radius-attribute**

Context config>aaa>l2tp-acct-plcy

Description This command enables the context to specify the RADIUS parameters that the system should include into RADIUS authentication-request messages.

The **no** form of this command disables the RADIUS attributes to be included in the policy.

nas-identifier

Syntax [**no**] **nas-identifier**

Context config>aaa>l2tp-acct-plcy>include-radius-attribute

Description This command enables the generation of the nas-identifier RADIUS attribute.

The **no** form of this command reverts to the default.

nas-port

Syntax [**no**] **nas-port** *bit-specification binary-spec*

Context config>aaa>l2tp-acct-plcy>include-radius-attribute

Description This command enables the generation of the nas-port RADIUS attribute. You enter decimal representation of a 32-bit string that indicates your port information. This 32-bit string can be compiled based on different information from the port (data types). By using syntax number-of-bits data-type you indicate how many bits from the 32 bits are used for the specific data type. These data types can be combined up to 32 bits in total. In between the different data types 0's and/or 1's as bits can be added.

The **no** form of this command disables the nas-port configuration.

Parameters *bit-specification binary-spec* — Specifies the NAS-Port attribute.

Values

binary-spec	<bit-specification> <binary-spec>
bit-specification	0 | 1 | <bit-origin>
bit-origin	*<number-of-bits><origin>
number-of-bits	1 to 32
origin	o | i | s | m | p

o outer VLAN ID
i inner VLAN ID
s slot number

m MDA number
p port number or lag-id

Output The following output shows an example.

Sample

```
*12o*12i00*2s*2m*2p => oooo oooo oooo iiii iiii iiii 00ss mmpp
If outer vlan = 0 & inner vlan = 1 & slot = 3 & mda = 1 & port = 1
=>  0000 0000 0000 0000 0000 0001 0011 0101 => nas-port = 309
```

nas-port-id

Syntax **nas-port-id**
nas-port-id [**prefix-string** *string*] [**suffix** *suffix-option*]
no nas-port-id

Context config>aaa>l2tp-acct-plcy>include-radius-attribute

Description This command enables the generation of the nas-port-id RADIUS attribute. Optionally, the value of this attribute (the SAP-id) can be prefixed by a fixed string and suffixed by the circuit-id or the remote-id of the client connection. If a suffix is configured, but no corresponding data is available, the suffix used is 0/0/0/0/0/0.

The **no** form of this command reverts to the default.

Parameters *string* — Specifies that a user configurable string be added to the RADIUS NAS port attribute, up to 8 characters.

suffix-option — Specifies the suffix type to be added to the RADIUS NAS port attribute.

Values circuit-id, remote-id

nas-port-type

Syntax **nas-port-type**
nas-port-type [*type*]
no nas-port-type

Context config>aaa>l2tp-acct-plcy>include-radius-attribute

Description This command enables the generation of the nas-port-type RADIUS attribute. If set to **nas-port-type**, the following is sent: values: 32 (null-encap), 33 (dot1q), 34 (qinq), 15 (DHCP hosts). The **nas-port-type** can also be set as a specified value, with an integer from 0 to 255.

The **no** form of this command reverts to the default.

Parameters *type* — Specifies an enumerated integer that specifies the value that is put in the RADIUS nas-port-type attribute.

Values 0 to 255

radius-accounting-server

Syntax **radius-accounting-server**

Context config>aaa>l2tp-acct-plcy>include-radius-attribute

Description This command enables the context for defining RADIUS accounting server attributes under a given session authentication policy.

access-algorithm

Syntax **access-algorithm** {**direct** | **round-robin**}
no access-algorithm

Context config>aaa>l2tp-acct-plcy>radius-acct-server

Description This command configures the algorithm used to access the list of configured RADIUS servers.

The **no** form of this command reverts to the default.

Default access-algorithm direct

Parameters **direct** — Specifies that the first server is used as primary server for all requests, the second as secondary and so on.

round-robin — Specifies that the first server is used as primary server for the first request, the second server as primary for the second request, and so on. If the router gets to the end of the list, it starts again with the first server.

retry

Syntax **retry** *count*

Context config>aaa>l2tp-acct-plcy>radius-acct-server

Description This command configures the number of times the router attempts to contact the RADIUS server for authentication.

Note: The retry count includes the first attempt.

The **no** form of this command reverts to the default value.

Default retry 3 (the initial attempt as well as two retried attempts)

Parameters *count* — Specifies the retry count.

Values 1 to 10

router

Syntax **router** *router-instance*
router service-name *service-name*
no router

Context config>aaa>l2tp-acct-plcy>radius-acct-server

Description This command specifies the number of times the router attempts to contact the RADIUS server for authentication, if not successful the first time.

The **no** form of this command reverts to the default value.

Parameters *router-instance* — Specifies the router instance.

Values

router-name | *vprn-svc-id*

router-name	Base, management Default - Base
vprn-svc-id	1 to 2147483647

service-name — Specifies the service name up to 64 characters.

server

Syntax **server** *server-index* **address** *ip-address* **secret** *key* [**hash** | **hash2**] [**port** *port*] [**create**]
no server *server-index*

Context config>aaa>l2tp-acct-plcy>radius-acct-server

Description This command adds a RADIUS server and configures the RADIUS server IP address, index, and key values.

Up to five RADIUS servers can be configured at any one time. RADIUS servers are accessed in order from lowest to highest index for authentication requests until a response from a server is received. A higher indexed server is only queried if no response is received from a lower indexed server (which implies that the server is not available). If a response from a server is received, no other RADIUS servers are queried.

The **no** form of this command removes the server from the configuration.

Parameters *server-index* — The index for the RADIUS server. The index determines the sequence in which the servers are queried for authentication requests. Servers are queried in order from lowest to highest index.

Values 1 to 16 (a maximum of 5 accounting servers)

address ip-address — The IP address of the RADIUS server. Two RADIUS servers cannot have the same IP address. An error message is generated if the server address is a duplicate.

secret *key* — The secret key to access the RADIUS server. This secret key must match the password on the RADIUS server.

Values secret-key — A string up to 20 characters.
hash-key — A string up to 33 characters.
hash2-key — A string up to 55 characters.

hash — Specifies the key is entered in an encrypted form. If the hash parameter is not used, the key is assumed to be in a non-encrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the hash parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form. If the hash2 parameter is not used, the less encrypted hash form is assumed.

port — Specifies the UDP port number on which to contact the RADIUS server for authentication.

Values 1 to 65535

source-address

Syntax **source-address** *ip-address*
no source-address

Context config>aaa>l2tp-acct-plcy>radius-acct-server

Description This command configures the source address of the RADIUS messages.

The **no** form of this command reverts to the default value.

Parameters *ip-address* — Specifies the source address to be used for NAT RADIUS accounting.

timeout

Syntax **timeout** [**sec** *seconds*] [**min** *minutes*]

Context config>aaa>l2tp-acct-plcy>radius-acct-server

Description This command configures the time that the router waits for a response from a RADIUS server.

The **no** form of this command reverts to the default value.

Default timeout sec 5

Parameters *seconds* — Specifies the time, in seconds, that the router waits for a response from a RADIUS server.

Values 1 to 59

minutes — Specifies the time, in minutes, that the router waits for a response from a RADIUS server.

Values 1 to 1

request-script-policy

Syntax **request-script-policy** *script-policy*
no request-script-policy

Context config>aaa>l2tp-acct-plcy>radius-acct-server

Description This command specifies the RADIUS script policy to be used for accounting-request packets.

The **no** form of this command removes the policy from the configuration.

Parameters *script-policy* — Configures a Python script policy name to modify Access-Request messages.

6.9.2.4 Show Commands

Note: The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

peer

Syntax **peer** *ip-address* [**udp-port** *port*]
peer *ip-address* statistics [**udp-port** *port*]
peer [**draining**] [**blacklisted** | **selectable** | **unreachable**]

Context show>router>l2tp

Description This command displays L2TP peer operational information/

Parameters *ip-address* — Specifies the IP dress for the L2TP peer.

Values

ip-address	ipv4-address - a.b.c.d
ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D
draining	keyword
statistics	keyword
port	[1 to 65535]

port — Specifies the local UDP port of this L2TP

Values 1 to 65535

draining — Specifies to display information about the L2TP peer being drained.

blacklisted — Specifies to display information about the peer has been deemed unreachable and has been put in the tunnel-selection-blacklist.

selectable — Specifies to display information about the peer has been deemed unreachable for the time specified in tmnxL2tpXtTuSelBlacklistMaxTime and can now be selected for a single session, after which the reachability of the peer is reevaluated.

unreachable — Specifies to display information about the sysUpTime when the this peer was deemed unreachable for the last time.

Output The following output is an example of L2TP peer information.

Sample Output

```
show router l2tp peer 10.100.0.2
===============================================================================
Peer IP: 10.100.0.2
===============================================================================
Roles capab/actual: LAC LNS /LAC  -     Draining          : false
Tunnels           : 1                   Tunnels Active    : 0
Sessions          : 1                   Sessions Active   : 0
Reachability      : blacklisted         Time Unreachable  : 01/31/2013 08:55:06
Time Blacklisted  : 01/31/2013 08:55:06 Remaining (s)     : 34
===============================================================================
Conn ID                     Loc-Tu-ID Rem-Tu-ID State              Ses Active
  Group                                                           Ses Total
    Assignment
-------------------------------------------------------------------------------
977207296                   14911     0         closed             0
  base_lac_base_lns                                               1
    t1
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
```

```
show router l2tp tunnel detail
===============================================================================
L2TP Tunnel Status
===============================================================================
Connection ID: 831782912
State        : closedByPeer
IP           : 10.0.0.1
Peer IP      : 10.100.0.2
Tx dst-IP    : 10.100.0.2
Rx src-IP    : 10.100.0.2
Name         : lac
Remote Name  :
Assignment ID: t1
Group Name   : base_lac_base_lns
Acct. Policy : l2tp-base
Error Message: N/A

                                       Remote Conn ID    : 4294901760
Tunnel ID         : 12692              Remote Tunnel ID  : 65535
UDP Port          : 1701               Remote UDP Port   : 1701
Preference        : 50                 Receive Window    : 64
Hello Interval (s): 300
Idle TO (s)       : 5                  Destruct TO (s)   : 60
Max Retr Estab    : 5                  Max Retr Not Estab: 5
Session Limit     : 32767              AVP Hiding        : sensitive
Transport Type    : udpIp              Challenge         : never
Time Started      : 01/31/2013 08:56:58 Time Idle        : 01/31/2013 08:56:58
Time Established  : N/A                Time Closed       : 01/31/2013 08:56:58
Stop CCN Result   : reqShutDown        General Error     : noError
Blacklist-state   : blacklisted
Blacklist Time    : 01/31/2013 08:56:58 Remaining (s)    : 49
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
```

l2tp

Syntax **l2tp**

Context show>system

Description This command displays L2TP system information.

Output The following output is an example of L2TP system information.

Sample Output

```
*A:Dut-C# show system l2tp
===============================================================================
L2TP system
===============================================================================
Non MC tunnel ID range                                 : 8193-16383
```

```
Max number of tunnels                                    : 16383
Max number of sessions                                   : 131071
Max number of sessions per tunnel                        : 32767
===============================================================================
```

sync

Syntax **sync** [**peer** *ip-address*] [**statistics**]
sync peer *ip-address* **detail**

Context show>redundancy>multi-chassis

Description This command displays synchronization information.

Parameters *ip-address* — Specifies the IP address of the peer.

Values ipv4-address - a.b.c.d

detail — Keyword to display detailed output.

statistics — Keyword to display statistics.

Output The following output is an example of multi-chassis sync information.

Sample Output

```
*A:Dut-C# show redundancy multi-chassis sync peer 10.1.2.2 detail
===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 10.1.2.2
Description             : Mc-Lag peer 10.1.2.2
Authentication          : Disabled
Source IP Address       : 10.1.1.1
Admin State             : Enabled
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     : SUBMGMT-PPPOE SRRP l2tp
Sync Admin State        : Up
Sync Oper State         : Up
Sync Oper Flags         :
DB Sync State           : inSync
Num Entries             : 2028
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
Rem Num Entries         : 2028
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
```

```
Rem OMCR Alarm Entries  : 0
===============================================================================
MCS Application Stats
===============================================================================
Application             : igmp
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : igmpSnooping
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : subMgmtIpoe
Num Entries             : 0
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : srrp
Num Entries             : 26
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 26
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : mcRing
Num Entries             : 0
```

```
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : mldSnooping
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : dhcpServer
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : subHostTrk
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : subMgmtPppoe
Num Entries             : 2000
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 2000
Rem Lcl Deleted Entries : 0
```

```
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : mcIpsec
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : mld
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : python
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : l2tp
Num Entries             : 2
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 2
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : diamProxy
Num Entries             : 0
Lcl Deleted Entries     : 0
```

```
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
===============================================================================
===============================================================================
Ports synced on peer 10.1.2.2
===============================================================================
Port/Encap                    Tag
-------------------------------------------------------------------------------
3/2/5
  1-999                       pppoe1
  1000-1000                   srrp1
3/2/6
  1-999                       pppoe2
===============================================================================
===============================================================================
DHCP Server instances synced on peer 10.1.2.2
===============================================================================
Router-Name                      Server-Name
  Tag
-------------------------------------------------------------------------------
No instances found
===============================================================================
===============================================================================
Python cache instances synced on peer 10.1.2.2
===============================================================================
Python-Policy                 Tag
-------------------------------------------------------------------------------
No instances found
===============================================================================
===============================================================================
L2TP instances
===============================================================================
Router         Tag                              SRRP
-------------------------------------------------------------------------------
Base           lac1                             1
Base           lac2                             2
===============================================================================
===============================================================================
Track SRRP instances
===============================================================================
SRRP                    : 1
-------------------------------------------------------------------------------
L2TP tunnel ID start    : 1
L2TP tunnel ID end      : 1

SRRP                    : 2
-------------------------------------------------------------------------------
L2TP tunnel ID start    : 2
L2TP tunnel ID end      : 2
===============================================================================
===============================================================================
```

```
Diameter proxy instances synced on peer 10.1.2.2
===============================================================================
Diameter-Peer-Policy            Tag
-------------------------------------------------------------------------------
No instances found
===============================================================================
===============================================================================
*A:Dut-C#
```

6.9.2.5 Debug Commands

assignment-id

Syntax **assignment-id** *assignment-id*

Context debug>router>l2tp

Description This command enables and configures debugging for the L2TP tunnel with a given assignment-id

Parameters *assignment-id* — Specifies a string that distinguishes this L2TP tunnel.

event

Syntax [**no**] **event**

Context debug>router>l2tp
debug>router>l2tp>assignment-id
debug>router>l2tp>group
debug>router>l2tp>peer
debug>router>l2tp>tunnel

Description This command configures an L2TP debugging event.

group

Syntax **group** *tunnel-group-name*

Context debug>router>l2tp

Description This command enables and configures debugging for an L2TP group.

Parameters *tunnel-group-name* — Specifies the tunnel group name up to 63 characters.

peer

Syntax **peer** *ip-address* [**udp-port** *port*]

Context debug>router>l2tp

Description This command enables and configures debugging for an L2TP peer.

Parameters *ip-address* — Specifies the IP address of the session.

Values ip-address : ipv4-address - a.b.c.d
ipv6-address - x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x - [0 to FFFF]H
d - [0 to 255]D

udp-port *port* — Specifies the local UDP port of this L2TP.

Values 1 to 65535

tunnel

Syntax **tunnel** *connection-id*

Context debug>router>l2tp

Description This command enables and configures debugging for an L2TP tunnel.

Parameters *connection-id* — Specifies the connection ID of the L2TP session associated with this session.

Values 1 to 4294967295

recovery

Syntax [**no**] **recovery**

Context debug>router>l2tp>assignment-id>event
debug>router>l2tp>event
debug>router>l2tp>group>event
debug>router>l2tp>peer>event
debug>router>l2tp>tunnel>event

Description This command configures L2TP LAC state recovery event debugging.

recovery-failed

Syntax [**no**] **recovery-failed**

Context debug>router>l2tp>assignment-id>event
debug>router>l2tp>event
debug>router>l2tp>group>event
debug>router>l2tp>peer>event
debug>router>l2tp>tunnel>event

Description This command configures L2TP LAC state recovery failed event debugging.

7 Triple Play Security

7.1 Triple Play Security Features

7.1.1 Anti-Spoofing Filters

This section describes the 7450 ESS or 7750 SR acting as a Broadband Subscriber Aggregator (BSA).

Anti-spoofing is useful to prevent certain packets from gaining unauthorized network access. Such packets originate from within a network and with an invalid source address.

Enabling anti-spoof filtering on a subscriber-facing SAP causes the anti-spoof table to be populated with all static and dynamic host information available on the SAP. Enabling anti-spoof filtering on the SAP will fail if any static hosts are defined without the proper addresses specified for the selected anti-spoof filter type.

When enabled, forwarding IP packets that ingress the SAP is dependent on a successful anti-spoof table match with an entry in the table. DHCP and non-IP packets (including ARP) are not subject to anti-spoof filtering. If an entry does not match the ingress packet, the packet is silently discarded while incrementing the SAP discard counter.

Anti-spoof filtering is only allowed on VPLS SAPs, IES SAP-based IP interfaces, and VPRN SAP-based IP interfaces. Anti-spoof filtering is not available on IES or VPRN SDP bound IP interfaces. Anti-spoof filtering is not supported on Epipe and other VLL type services. Support for anti-spoofing is dependent on SAP based service interfaces. Note VPRN and VLL are supported on the 7750 SR only.

Note: Anti-spoofing filters, with type **ip-mac**, must be enabled to perform Enhanced Subscriber Management (as described in the Triple Play Enhanced Subscriber Management section).

7.1.1.1 Anti-spoofing Filter Types

A SAP or interface that supports anti-spoof filtering can be configured to use one of three types of anti-spoof tables. The type of table used by the SAP is dependent on the type of anti-spoof filtering desired, only one anti-spoofing table type is supported per SAP:

- When only the incoming source MAC address is to be verified, the source MAC table must be defined (anti-spoof type **mac**).
- When only the incoming source IP address is to be verified, the source IP table must be defined (anti-spoof type **ip**).
- When both the incoming source MAC and source IP addresses are to be verified, the combination source IP and source MAC table must be defined (anti-spoof type **ip-mac**).

Note: Setting the anti-spoof filter type for the SAP is dependent on pre-existing static host definitions, for example, attempting to set the SAP anti-spoof filtering to **mac** will fail if any static hosts exist that do not have a defined MAC address.

The anti-spoof table of a SAP or interface is populated from the DHCP lease state table and from any statically defined hosts on the SAP or interface.

7.1.1.2 Filtering Packets

Packets from a client that match an anti-spoof filter entry when anti-spoof filtering is enabled are allowed to be further processed by the system. The matching packet is still subject to other forwarding criteria including potentially ACL filtering.

All packets that are not exempt from anti-spoofing and do not match an entry in the anti-spoof table are discarded. Every discard event will increment the SAP discard packet counter. The discard event is not logged or alarmed, but a threshold alarm could be configured for the counter (see Configuring System Monitoring Thresholds in the 7450 ESS, 7750 SR, 7950 XRS, and VSR Basic System Configuration Guide).

Not all ingress packets are subject to the anti-spoof filtering when enabled. Non-IP packets are exempt for anti-spoof filter lookups and are allowed to be further processed by the system. This includes ARP requests and replies, as well as PPPoE packets. The only IP packets exempt from anti-spoof filtering are DHCP packets destined to the server UDP port 67. DHCP packets destined to the client UDP port number (port 68) are not exempt.

7.1.2 Layer 2 Triple Play Security Features

7.1.2.1 MAC Pinning

This section describes the 7450 ESS and the 7750 SR acting as a Broadband Subscriber Aggregator (BSA).

DHCP snooping and IP and MAC auto-filters can be used to prevent Theft of Service (by a malicious user spoofing another user's address). However, these auto-filters do not discard non-IP packets such as PPPoE packets, thus potentially allowing a MAC address to be relearned on another SAP. MAC pinning closes this loophole, by not allowing a MAC address to be relearned on another SAP.

When MAC pinning is enabled, a MAC address learned on one SAP or SDP can not be relearned on another SAP or SDP in the same VPLS, until the FDB entry for the MAC address times out. (Note that in case MAC aging is disabled, MAC entries on a SAP or SDP with MAC pinning enabled will effectively become permanent.)

MAC pinning is implicitly enabled when DHCP auto-filters are enabled, and cannot be disabled. For MAC addressing learned during DHCP address assignment (when DHCP snooping function is active at least on one port of the VPLS), the MAC address is tied to a given SAP for the duration of the DHCP lease.

7.1.2.2 MAC Protection

In a Layer 2 environment, a malicious subscriber could create a denial-of-service (DoS) attack by sending Ethernet frames, with as source MAC address the address of a gateway (for example, the IP next hop upstream). As MAC learning is typically enabled, this would move the learned gateway MAC from the uplink SAP or SDP to the subscriber’s SAP, causing all communication to the gateway to be disrupted. If a local content server is attached to the same VPLS, a similar attack could be launched against it.

Communication between subscribers can be disallowed using Split Horizon Groups, but this by itself is not sufficient to prevent such an attack. The solution is to create a mechanism to explicitly protect some MAC addresses against being relearned on other SAPs.

The **mac-protect** feature on the 7450 ESS and 7750 SR allows a list of special MAC addresses to be configured in a VPLS. Two checks can then be made on incoming packets against these protected MAC addresses:

- **[no] auto-learn-mac-protect**: Used to enable the automatic protection of source MAC addresses learned on the associated object. MAC protection is used in conjunction with **restrict-protected-src**, **restrict-unprotected-dst** and **mac-protect**. When this command is applied or removed, the MAC addresses are cleared from the related object.
- **restrict-protected-src**: Used to prevent DoS attacks. If the source MAC address of a packet from a subscriber matches a protected entry, it is probable that this subscriber tried to impersonate the gateway or server. If no parameter is specified, such packets are discarded, a trap is generated, and the SAP on which it arrived is placed operationally down. If the **alarm-only** parameter is specified, the packet is forwarded, and an alarm is generated but the source MAC is not learned. If the **discard-frame** parameter is specified, the packet is discarded and an alarm generated.
- **restrict-unprotected-dst**: Used to force traffic from subscribers to only go towards a few defined destinations (the gateways or servers). Any packet from a subscriber whose destination MAC address does not match a protected entry is discarded.

7.1.2.3 DoS Protection

This section describes the mechanisms and limitations of DoS protection related to subscriber management snooping functions. This feature is only supported on 7750 SR-Series and 7450 ESS-Series redundant chassis models. In subscriber aggregation networks, these routers play an active role in several protocols. Subscribers either intentionally or unknowingly interfere with the operation of the node’s processing capacity (for example, excessive ARP handling) or other user traffic.

Routing protocols such as OSPF and ISIS could also be a threat as packets can be injected by customers (erroneously or maliciously) which could cause high CPM overload. Service providers are concerned about DoS protection including DoS attacks when acting as a subscriber aggregation device and guarding against DoS attacks using unprovisioned protocols.

7.1.2.3.1 Subscriber Aggregation Network

In a subscriber aggregation network, multiple devices such as the 7750 SR or 7450 ESS routers provide access to a DHCP or a RADIUS server. These servers usually do not scale high enough to provide the means to control access to snooping functions through a controlled queue. It is possible, under severe conditions, that the network could become unavailable if the node cannot handle requests from subscribers.

Because the IOMs cannot be scaled to provide a per-subscriber queue to control traffic, a monitoring function, handled by the CPM, is provided. With this monitoring system, the CPM tracks the number of control plane messages set per subscriber and limits the rate to a specified level and provides feedback using event generation to alert a centralized system of a possible DoS attack.

The CPM provides a prioritized access to the CPU. Since the number of control packets expected from a subscriber should have a low rate, and under normal conditions, the system will provide a rate limit on a per subscriber or MAC basis and will drop a subscriber control packet before it is queued or processed by the CPU. The system is configured with expected arrival rate of per MAC or subscriber control packet rates and optionally total rate per interface/SAP.

The system maintains a per-second running rate monitor per SAP and per MAC. If an entry is using more than the configured rate, the system will not forward that packet to be queued. Every existing subscriber host is monitored. A subscriber host is flagged and the system observed with an excessive rate of control packets. With PPPoE, the CPM monitors subscriber hosts before the IP address is provided by the SAP, MAC, or session-id combination.

The control protocols affected by this mechanism include:

- ARP (in arp-reply-agent)
- DHCP (for discover and renew)
- ICMP
- PPPoE
- IGMP

7.1.2.3.2 Network Control Filtering

The 7750 SR or 7450 ESS can block network control traffic for unconfigured protocols. For example, if OSPF is not configured on an IP interface, all OSPF-related traffic should be dropped before the traffic reaches the CPU.

Protocols are blocked based on whether that particular protocol is configured to run on the given IP interface. It is not required to re-configure the permitted protocols.

Protocol traffic control by this mechanism includes:

- OSPFv2
- OSPFv3
- IS-IS
- RSVP-TE
- LDP
- RIP
- PIM
- MLD
- IGMP
- BGP
- BFD
- L2PT
- PPP
- DHCP

7.1.2.4 VPLS Redirect Policy

This section describes the 7450 ESS or 7750 SR acting as a Broadband Subscriber Aggregator (BSA) with Layer 2 aggregation towards a Broadband Subscriber Router (BSR).

In a Triple Play network it may be desired to route some traffic from/to subscribers through a Deep Packet Inspection (DPI) device, for example, to limit peer-to-peer traffic. However such a DPI device typically has limited bandwidth available, so only those packets that need inspection should be sent to it.

In a Layer 3 network, such policy-based redirection can be achieved using “next-hop redirect” ACL entries. In a layer 2 (VPLS) aggregation network, the same result can be achieved using “redirect to SAP” or “redirect to SDP” policy.

Refer to the ACL Next-Hop for VPLS section in the 7450 ESS, 7750 SR, 7950 XRS, and VSR Services Overview Guide.

7.1.3 ARP Handling

7.1.3.1 ARP Reply Agent

This section describes the 7450 ESS or 7750 SR acting as a Broadband Subscriber Aggregator (BSA).

In Triple Play networks, typically downstream broadcast is not allowed on subscriber SAPs. As a result, subscribers can not receive ARP requests from the network. Instead, the 7450 ESS or 7750 SR responds to ARP requests from the network, with information from the DHCP lease state table.

In the upstream direction (towards the network), the ARP reply agent intercepts ARP Requests on subscriber SAPs, and checks them against the DHCP lease state table. The purpose is to prevent a malicious subscriber spoofing ARP Request or ARP reply messages and thus populating the upstream router's ARP table with incorrect entries.

When the keyword **sub-ident** is added in the ARP reply agent configuration, also the subscriber identity is checked. If an upstream ARP request is targeted to the same subscriber, it is dropped. Otherwise, it is flooded to all VPLS interfaces outside the received Split Horizon Group (SHG).

Static hosts can be defined on the SAP using the **host** command. Dynamic hosts are enabled on the system by enabling the **lease-populate** command in the SAP's **dhcp** context. In the event that both a static host and a dynamic host share the same IP and MAC address, the VPLS ARP reply agent retains the host information until both the static and dynamic information are removed. In the event that both a static and dynamic host share the same IP address, but different MAC addresses, the VPLS ARP reply agent is populated with the static host information.

In brief, the ARP Replay Agent operation is as follows:

- For ARP request received from a customer SAP:
  - first check in DHCP lease state table - if no match: discard
  - if (**sub-ident** enabled) and (destination = same subscriber): discard
  - otherwise: flood to all SAPs/SDPs outside this SHG
- For ARP request received from the network:
  - lookup IP address in DHCP lease state table - if no match: discard
  - otherwise: respond with MAC address from the DHCP lease state table

7.1.3.2 Dynamic ARP Table Population

This section describes the 7450 ESS or 7750 SR acting as a Broadband Subscriber Aggregator (BSA) with Layer 3 forwarding towards the network.

In an IES service, the system's ARP table can be populated dynamically using entries in the DHCP lease state table (in turn populated from snooping DHCP ACK messages (see DHCP Snooping)), and from static hosts defined on the SAP. In the router ARP table these are indicated with state managed.

In the event that both a static host is created with the same IP and MAC address as an existing managed entry, creation will fail and a trap is generated.

In the event that a DHCP Lease needs to be populated with the same IP and MAC address as an existing static host entry, creation will fail and a trap is generated.

No **static-arp** creation is possible when combined with **arp-populate**.

7.1.3.3 Local Proxy ARP

This section describes the 7450 ESS or 7750 SR acting as a Broadband Subscriber Aggregator (BSA) with Layer 3 forwarding towards the network.

Local proxy ARP allows the 7450 ESS or 7750 SR to respond to ARP requests received on an interface, for an IP address which is part of a subnet assigned to the interface. When the local proxy ARP feature is enabled, the switch responds to all ARP requests for IP addresses belonging to the subnet with the MAC address of the interface, and forwards all traffic between hosts in the subnet.

This feature is intended to be used in situations (such as DSL aggregation networks) when hosts belonging to the same subnet are prevented from directly communicating with each other over the subnet by the configuration of the switch (or DSLAM) to which they are connected.

Note: When **local-proxy-arp** is enabled under a IES service, all ICMP redirects on the ports associated with the service are automatically blocked. This prevents users from learning each other's MAC address (from ICMP redirects).

The implementation of proxy ARP with support for local proxy ARP allows the router to respond to ARP requests in the subnet assigned to an IES or VPRN interface, thus allowing multiple customers to share the same IP subnet.

7.1.4 Web Portal Redirect

This section describes the 7450 ESS or 7750 SR acting as a Broadband Subscriber Aggregator (BSA).

In a Triple Play network it may not be desired (or feasible) to perform manual provisioning of new services and service changes. The ideal way of working is automatic provisioning, with the end-user supplying his details at a retailer, or (if physical connectivity is already present through an on-line customer portal).

The 7450 ESS and 7750 SR support a special ACL that automatically redirects subscribers to a predefined URL. This is done by sending a HTTP 302 (moved) message to the subscriber, regardless of the requested URL.

The message flow is as follows (see Figure 48 below):

a. The subscriber gets an IP address using DHCP (if the customer is trying to use a static IP he is blocked by the anti-spoofing filter).
b. The subscriber tries to connect to a website (TCP SYN, TCP ACK, HTTP GET).
c. The 7450 ESS or 7750 SR intercepts the HTTP GET request and discards it.
d. The 7450 ESS or 7750 SR then responds to the subscriber with a HTTP 302 message (service temporarily unavailable or moved), containing a new target URL (that of the portal) configured by the operator. This target URL can include the subscriber's IP and MAC addresses as part of the portal's URL string.
e. The subscriber's web browser closes the original TCP connection and opens a new connection to the web portal where the subscriber can sign up or change his/her service Profile.
f. After approving the changes, the web portal updates the ACL (directly or through another system such as the Nokia 5750 SSC) to remove the redirection policy.
g. The subscriber can now connect to the original site.

Figure 48 **IP Illustration of Message Flow in Web Portal Redirect**

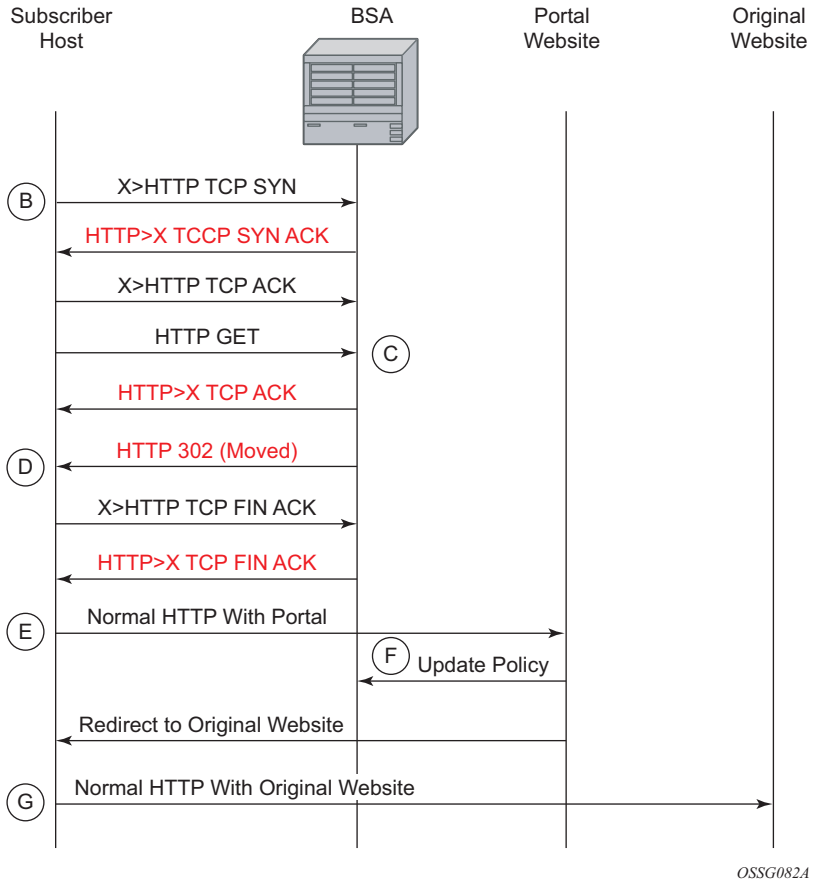

The items in red text in Figure 48 are messages the 7450 ESS or 7750 SR sends (masquerading as the destination), regardless of the destination IP address or type of service.

The following displays information that can optionally be added as variables in the portal URL (http-redirect url):

- $IP – Customer's IP address
- $MAC – Customer's MAC address
- $URL – Original requested URL
- $SAP – Customer's SAP
- $SUB – Customer's subscriber identification string
- $CID – The string that represents the circuit-id or interface-id of the subscriber host (hexadecimal format)

- $RID – The string that represents the remote-id of the subscriber host (hexadecimal format)

Note: The subscriber's IP and MAC address variables are populated from the anti-spoofing list, and thus anti-spoofing must be enabled (see section Anti-Spoofing Filters).

Since most websites are accessed using the domain name, the 7450 ESS or 7750 SR needs to allow DNS queries, and an ACL entry to this effect should be included in the filter (see Configuring Web Portal Redirect).

7.2 Configuring Triple Play Security with CLI

7.2.1 Common Configuration Tasks

7.2.1.1 Configuring Anti-Spoofing Filters

Anti-spoofing filters are used to prevent malicious subscribers from sending IP packets with a forged IP or MAC address, and thus mis-directing traffic. The anti-spoofing filter is populated from the DHCP lease state table, and DHCP snooping must be enabled on the SAP.

There are three types of filters (MAC, IP, and IP+MAC). One type is allowed per SAP.

The following displays an IES service interface configuration with anti-spoofing.

```
A:ALA-48>config>service>ies# info
----------------------------------------------
            interface "test123" create
                address 10.10.42.41/24
                local-proxy-arp
                proxy-arp
                    policy-statement "ProxyARP"
                exit
                sap 1/1/7:0 create
                    anti-spoof ip
                exit
                arp-populate
                dhcp
                    lease-populate 1
                    no shutdown
                exit
            exit
            no shutdown
----------------------------------------------
A:ALA-48>config>service>ies#
```

7.2.1.2 Configuring Triple Play Security Features

7.2.1.2.1 Configuring MAC Pinning

The following example displays a partial BSA configuration with MAC pinning enabled on a SAP.

```
A:ALA-48>config>service# info
----------------------------------------------
vpls 800 customer 6001 create
    description "VPLS with residential split horizon for DSL"
    stp
        shutdown
    exit
    sap 2/1/4:100 split-horizon-group "DSL-group2" create
        description "SAP for RSHG"
              mac-pinning
    exit
    no shutdown
----------------------------------------------
A:ALA-48>config>service#
```

7.2.1.2.2 Configuring MAC Protection

Preventing Access By Residential Subscribers Using Protected (Gateway) MAC Addresses

The first step is to create a list of MAC addresses to be protected. The second step is to prevent access using these source addresses inside an SHG or a SAP.

The following example displays a partial BSA configuration with some protected MAC addresses on any SAP created inside the SHG.

```
A:ALA-48>config>service# info
----------------------------------------------
vpls 800 customer 6001 create
    no shutdown
    split-horizon-group "mygroup" create
        restrict-protected-src
    exit
    description "VPLS with residential split horizon for DSL"
    mac-protect
        mac 00:00:17:FE:82:D8
        mac 93:33:00:00:BF:92
    exit
----------------------------------------------
A:ALA-48>config>service#
```

Restricting Access By Residential Subscribers To a Small List Of Upstream MAC Addresses

The first step is to create a list of MAC addresses to be protected. The second step is to restrict access to these addresses only from an SHG or a SAP. (If the MAC address of an upstream server is not known, it can be discovered using, for example, the CPE ping OAM tool.)

The following example displays a partial BSA configuration with restricted access to some MAC addresses from a specified SAP (an unrestricted access from any other SAP within the VPLS).

```
A:ALA-48>config>service# info
----------------------------------------------
vpls 800 customer 6001 create
    no shutdown
    description "VPLS with restricted access on a SAP"
    mac-protect
        mac 00:00:17:FE:82:D8
        mac 93:33:00:00:BF:92
    exit
    sap 1/1/4:30 create
        restrict-unprotected-dst
    exit
----------------------------------------------
A:ALA-48>config>service#
```

7.2.1.2.3 Configuring VPLS Redirect Policy

Figure 49 displays an IP filter entry configuration for VPLS redirect policy.

Figure 49 **VPLS Redirect Policy Example**

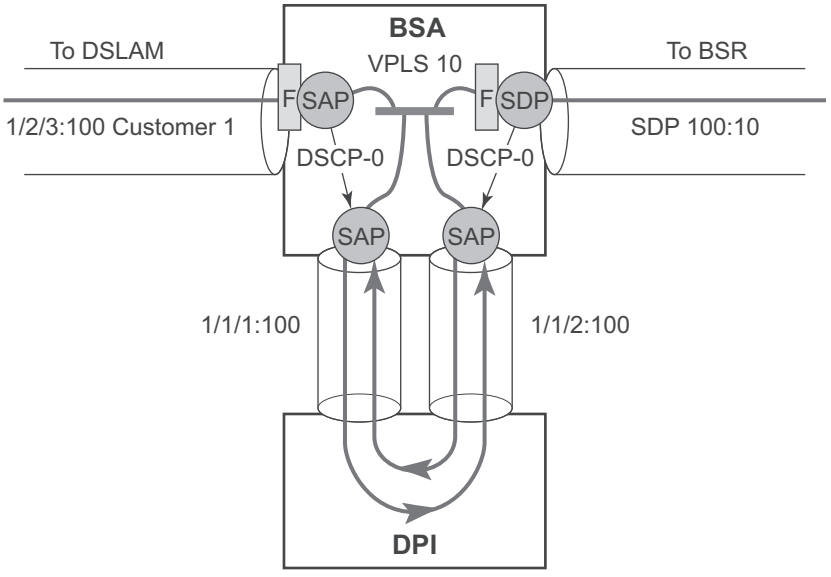

Information about defining and applying IP and MAC filters is described in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Router Configuration Guide*.

Creating the Filter

The following displays a redirect filter entry:

```
A:ALA-A>config>filter# info
----------------------------------------------
    ip-filter 10
        default-action forward
        entry 10
            match
                dscp be
            exit
            action forward next-hop sap 1/1/1:100
            exit
        exit
    exit
    ip-filter 11
        default-action forward
        entry 10
            match
                dscp be
            exit
                dscp be
            exit
            action forward next-hop sap 1/1/2:100
        exit
    exit
----------------------------------------------
A:ALA-A>config>filter#
```

Applying the Filter to a VPLS Service

The following displays how the redirection filter configured above is assigned to the ingress SAP from the DSLAM, and the ingress SDP from the BSR:

```
A:ALA-A>config>service>vpls# info
----------------------------------------------
    vpls 10 customer 1 create
        description "vpls10"
            sap 1/2/3:100 create
                ingress ip filter 10
            exit
            sap 1/1/1:100 create
            exit
            sap 1/1/2:100 create
            exit
            mesh-sdp 100:10 create
                ingress ip filter 11
            exit
        exit
    exit
----------------------------------------------
A:ALA-A>config>service>vpls#
```

7.2.1.3 Configuring ARP Handling

7.2.1.3.1 Configuring Proxy ARP

The implementation of proxy ARP with support for local proxy ARP allows the 7450 ESS or 7750 SR to respond to ARP requests in the subnet assigned to an IES or VPRN interface.

Configuring this command allows multiple customers to share the same IP subnet.

The following example displays an IES proxy ARP configuration:

```
A:ALA-48>config>service>ies# info
----------------------------------------------
            interface "test123" create
                address 10.10.42.41/24
                local-proxy-arp
                proxy-arp-policy "ProxyARP"
                exit
            exit
            no shutdown
----------------------------------------------
A:ALA-48>config>service>ies#
```

7.2.1.3.2 Configuring Local Proxy ARP

When local proxy ARP is enabled on an IP interface, the 7450 ESS or 7750 SR responds to all ARP requests for IP addresses belonging to the subnet with it's own MAC address, and forwards all traffic between hosts in that subnet. Local proxy ARP is disabled by default.

Note: When **local-proxy-arp** is enabled under a IES or VPRN service, all ICMP redirects on the ports associated with the service are automatically blocked. This prevents users from learning each other's MAC address (from ICMP redirects).

The following example displays a local proxy ARP IES configuration:

```
A:ALA-A>config>service>ies# info
----------------------------------------------
        interface "test" create
            shutdown
            address 10.10.36.2/24
            local-proxy-arp
        exit
----------------------------------------------
A:ALA-A>config>service>ies#
```

7.2.1.3.3 Configuring ARP Reply Agent in a VPLS Service

When ARP reply agent is enabled, the 7450 ESS or 7750 SR responds to ARP requests from the network, with information from the DHCP lease state table.

In the upstream direction (toward the network), the ARP reply agent intercepts ARP requests on subscriber SAPs, and checks them against the DHCP lease state table. The purpose is to prevent a malicious subscriber spoofing ARP request or ARP reply messages and thus populating the upstream router's ARP table with incorrect entries.

The following example displays a partial BSA configuration with ARP Reply Agent enabled on a SAP:

```
A:ALA-48>config>service# info
----------------------------------------------
...
    vpls 800 customer 6001 create
        description "VPLS with ARP Reply Agent active"
        sap 2/1/4:100 split-horizon-group "DSL-group2" create
    arp-reply-agent sub-ident
        exit
        sap 3/1/4:200 split-horizon-group "DSL-group2" create
    arp-reply-agent sub-ident
        exit
        no shutdown
...
----------------------------------------------
A:ALA-48>config>service#
```

7.2.1.3.4 Configuring Remote Proxy ARP

The following example displays the IES configuration to enable remote proxy ARP:

```
A:ALA-49>config>service>ies# info
----------------------------------------------
            interface "test-1A" create
                address 10.10.26.3/24
                remote-proxy-arp
            exit
            no shutdown
----------------------------------------------
A:ALA-49>config>service>ies#
```

7.2.1.3.5 Configuring Automatic ARP Table Population in an IES or VPRN Interface

The following example displays the IES DHCP configuration to enable automatic population of the ARP table using snooped DHCP information on an IES or VPRN (VPRN is supported on the 7750 SR only) interface:

```
A:ALA-1>config>service>ies>if# info
----------------------------------------------
                arp-populate
                dhcp
                    description "snooping_only"
                    lease-populate 1
                    no shutdown
                exit
----------------------------------------------
A:ALA-1>config>service>ies>if#


A:ALA-1>config>service>vprn>if# info
----------------------------------------------
                dhcp
                    description "test"
                    lease-populate 1
                    no shutdown
                exit
----------------------------------------------
A:ALA-1>config>service>ies>if#
```

7.2.1.3.6 Configuring CPU Protection

CPU Protection can be used to protect the SR OS in subscriber management scenarios. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR System Management Guide* for information about CPU Protection operation and configuration.

7.2.1.4 Configuring Web Portal Redirect

The generic CLI structure for defining and applying IP and MAC filters is described in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Router Configuration Guide*.

The following example displays an IP filter entry configuration for web-portal redirect:

```
A:ALA-A>config>filter# info
----------------------------------------------
    ip-filter 10 create
        description "filter to forward DNS and web traffic to my portal; redirect al
l other web traffic to the portal and drop everything else"
        default-action drop
```

```
            entry 10 create
                description "allows DNS traffic"
                    match protocol 17
                    dst-port 53
                exit
                action forward
            exit
            entry 20 create
                description "allows web traffic destined to portal (IP address 10.0.0.1)
"
                match protocol 6
                    dst-port eq 80
                    dst-ip 10.0.0.1
                exit
                action forward
            exit
            entry 30
                description "redirects all web traffic to portal"
                match protocol 6
                    dst-port eq 80
                exit
                action http-redirect http://www.myportal.com/defaultportal/
login.cgi?ip=$IP&mac=$MAC&orig_url=$URL&usb=$SUB
            exit
        exit
-------------------------------------------------
A:ALA-A>config>filter#
```

- Filter entry 10 in the example output allows the customer to access DNS to get the IP address of the original website they are trying to view.
- Entry 20 allows HTTP packets destined to the captive portal itself to be forwarded.

Note: The actual IP address (a.b.c.d) must be entered, not the DNS name ("www.myportal.com"). The IP address can be easily resolved from the 7450 ESS or 7750 SR CLI using the **ping** command.

- Entry 30 (which is the last option that does not drop the customer packets) checks for HTTP protocol and then starts the redirection process:
  - The 7450 ESS or 7750 SR intercepts the HTTP GET from the subscriber and respond with an HTTP 302 (temporarily moved) with the URL configured in the filter entry. This URL can contain some variables, notably the customer IP and MAC addresses to allow the portal to create an entry for the customer. The original requested URL is also included to redirect the client site back to the original requested site when the process is done.
  - The client then closes the connection with the original IP address and open a connection to the redirected server. Entry 20 allows this connection.

The following displays how the redirection filter configured above is assigned to an ingress SAP:

```
A:ALA-A>config>service>vpls# info
----------------------------------------------
    vpls 3 customer 6 create
        description "VPLS with web portal redirection filter applied"
        sap 2/1/5:0 create
            ingress
                filter ip 10
            exit
        exit
        no shutdown
    exit
----------------------------------------------
A:ALA-A>config>service>vpls#
```

7.3 Triple Play Security Command Reference

7.3.1 Command Hierarchies

- Anti-Spoofing Commands
- Layer 2 Security Commands
- ARP Handling Commands
- Show Commands
- Tools Commands

Note: This command tree is limited to those commands specific to Triple Play security. For the full command trees of a particular service type, consult the Services Overview Guide.

7.3.1.1 Anti-Spoofing Commands

```
config
    — service
        — vpls service-id [customer customer-id]  [vpn vpn-id] [m-vpls]
            — sap sap-id [split-horizon-group group-name]
                — anti-spoof {ip | mac | ip-mac}
                — no anti-spoof

config
    — service
        — ies service-id [customer customer-id] [vpn vpn-id]
            — [no] interface ip-int-name
                — sap sap-id
                    — anti-spoof {ip | mac | ip-mac}
                    — no anti-spoof
            — subscriber-interface
                — group-interface ip-int-name
                    — sap sap-id
                        — anti-spoof {ip | ip-mac}
                        — no anti-spoof
```

7.3.1.2 Layer 2 Security Commands

```
config
    — service
```

```
— [no] vpls service-id [customer customer-id]  [vpn vpn-id] [m-vpls]
    — mac-protect
        — [no] mac [ieee-mac-address]
    — [no] split-horizon-group [group-name] [residential-group]
        — [no] restrict-protected-src [alarm-only | discard-frame]
        — [no] restrict-unprotected-dst
    — mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan}]
        — [no] mac-pinning
        — [no] restrict-protected-src [alarm-only | discard-frame]
        — [no] restrict-unprotected-dst
    — spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan | vpls}] [split-horizon-group
      group-name]
        — [no] mac-pinning
        — [no] restrict-protected-src [alarm-only | discard-frame]
        — [no] restrict-unprotected-dst
```

7.3.1.3 ARP Handling Commands

```
config
    — service
        — [no] vpls service-id [customer customer-id]  [vpn vpn-id] [m-vpls]
            — sap sap-id
                — arp-reply-agent [sub-ident]
                — no arp-reply-agent

config
    — service
        — ies service-id [customer customer-id] [vpn vpn-id]
            — [no] interface ip-int-name
                — [no] arp-populate
                — arp-timeout seconds
                — no arp-timeout
                — [no] local-proxy-arp
                — proxy-arp-policy policy-name [policy-name]
                — no proxy-arp-policy
                — [no] remote-proxy-arp
                — static-arp ip-address ieee-address
                — no static-arp ip-address [ieee-address]
```

7.3.1.4 Show Commands

```
show
    — service
        — id service-id
            — arp [ip-address]|[mac ieee-address]|[sap sap-id]|[interface ip-int-name]|[sdp
              sdp-id:vc-id] | [summary]
```

7.3.1.5 Tools Commands

```
tools
    — perform
        — security
            — authentication-server-check server-address ip-address [port port] {{user-
              name user-name password password} | attr-from-file file-url} secret key
              [source-address ip-address] [timeout seconds] [router router-instance |
              service-name service-name] [debug]
```

7.3.2 Command Descriptions

7.3.2.1 Triple Play Security Configuration Commands

7.3.2.1.1 Triple Play Anti-Spoofing Commands

anti-spoof

Syntax **anti-spoof** {**ip** | **mac** | **ip-mac**}
no anti-spoof

Context config>service>ies>if>sap
config>service>vpls>sap
config>service>vprn>if>sap

Description This command enables anti-spoof filtering and optionally changes the anti-spoof matching type for the SAP.

The type of anti-spoof filtering defines what information in the incoming packet is used to generate the criteria to lookup an entry in the anti-spoof filter table. The type parameter (**ip**, **mac**, **ip-mac**) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled.

The **no** form of the command disables anti-spoof filtering on the SAP.

Default no anti-spoof

Parameters **ip** — Configures SAP anti-spoof filtering to use only the source IP address in its lookup. If a static host exists on the SAP without an IP address specified, the **anti-spoof ip** command will fail.

mac — Configures SAP anti-spoof filtering to use only the source MAC address in its lookup. If a static host exists on the SAP without a specified MAC address, the **anti-spoof mac** command will fail.

ip-mac — Configures SAP anti-spoof filtering to use both the source IP address and the source MAC address in its lookup.

anti-spoof

Syntax **anti-spoof** {**ip** | **ip-mac**}
no anti-spoof

Context config>service>ies>sub-if>grp-if>sap

Description This command enables anti-spoof filtering and optionally change the anti-spoof matching type for the SAP.

The type of anti-spoof filtering defines what information in the incoming packet is used to generate the criteria to lookup an entry in the anti-spoof filter table. The type parameter (**ip**, **ip-mac**) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled.

The **no** form of the command disables anti-spoof filtering on the SAP.

Default no anti-spoof

Parameters **ip** — Configures SAP anti-spoof filtering to use only the source IP address in its lookup. If a static host exists on the SAP without an IP address specified, the **anti-spoof ip** command will fail.

ip-mac — Configures SAP anti-spoof filtering to use both the source IP address and the source MAC address in its lookup.

7.3.2.1.2 Triple Play Layer 2 Security Commands

split-horizon-group

Syntax [**no**] **split-horizon-group** [*group-name*] [*residential-group*]

Context config>service>vpls

Description This command creates a new split horizon group for the VPLS instance. Traffic arriving on a SAP or spoke SDP within this split horizon group will not be copied to other SAPs or spoke SDPs in the same split horizon group.

A split horizon group must be created before SAPs and spoke SDPs can be assigned to the group.

The split horizon group is defined within the context of a single VPLS. The same group-name can be re-used in different VPLS instances.

Up to 30 split horizon groups can be defined per VPLS instance.

The **no** form of the command removes the group name from the configuration.

Default none

Parameters *group-name* — Specifies the name of the split horizon group to which the SDP belongs.

residential-group — Defines a split horizon group as a residential split horizon group (RSHG). Doing so entails that:

a. SAPs which are members of this Residential Split Horizon Group will have:
   - Double-pass queuing at ingress as default setting (can be disabled)
   - STP disabled (can not be enabled)
   - ARP reply agent enabled per default (can be disabled)
   - MAC pinning enabled per default (can be disabled)
   - Besides the multicast downstream also broadcast packets are discarded thus also blocking the unknown, flooded traffic

b. Spoke SDPs which are members of this Residential Split Horizon Group will have:
   - Downstream multicast traffic supported
   - Double-pass queuing is not applicable
   - STP is disabled (can be enabled)
   - ARP reply agent is not applicable (dhcp-lease-states are not supported on spoke SDPs)
   - MAC pinning enabled per default (can be disabled)

mac-protect

Syntax [**no**] **mac-protect**

Context config>service>vpls

Description This command indicates if this MAC is protected on the MAC protect list. When enabled, the agent will protect the MAC from being learned or re-learned on a SAP, spoke SDP or mesh SDP that has restricted learning enabled. The MAC protect list is used in conjunction with **restrict-protected-src**, **restrict-unprotected-dst** and **auto-learn-mac-protect**.

The **no** form of the command reverts to the default.

Default none

mac

Syntax [**no**] **mac** *ieee-address*

Context config>service>vpls>mac-protect

Description This command specifies the 48-bit IEEE 802.3 MAC address.

The **no** form of the command reverts to the default.

Default none

Parameters *ieee-address* — Specifies the 48-bit MAC address in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers.

auto-learn-mac-protect

Syntax [**no**] **auto-learn-mac-protect**

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls >mesh-sdp
config>service>vpls>split-horizon-group
config>service>vpls>endpoint
config>service>pw-template
config>service>pw-template>split-horizon-group

Description This command enables the automatic protection of source MAC addresses learned on the associated object. MAC protection is used in conjunction with restrict-protected-src, restrict-unprotected-dst and mac-protect. When this command is applied or removed, the MAC addresses are cleared from the related object.

When the auto-learn-mac-protect is enabled on an SHG the action only applies to the associated SAPs (no action is taken by default for spoke SDPs in the SHG). In order to enable this function for spoke SDPs within a SHG, the auto-learn-mac-protect must be enabled explicitly under the spoke SDP. If required, auto-learn-mac-protect can also be enabled explicitly under specific SAPs within the SHG.

The **no** form of the command reverts to the default.

Default no auto-learn-mac-protect

restrict-protected-src

Syntax **restrict-protected-src** [**alarm-only** | **discard-frame**]
no restrict-protected-src

Context config>service>vpls>sap

config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp
config>service>vpls>split-horizon-group
config>service>vpls>endpoint
config>service>pw-template>
config>service>pw-template>split-horizon-group

Description This command indicates how the agent will handle relearn requests for protected MAC addresses, either manually added using the mac-protect command or automatically added using the **auto-learn-mac-protect** command. While enabled all packets entering the configured SAP, spoke SDP, mesh SDP, or any SAP that is part of the configured split horizon group (SHG) is verified not to contain a protected source MAC address. If the packet is found to contain such an address, the action taken depends on the parameter specified on the **restrict-protected-src** command, namely:

- No parameter — The packet is discarded, an alarm is generated and the SAP, spoke SDP or mesh SDP is set operationally down. The SAP, spoke SDP or mesh SDP must be shutdown and enabled (no shutdown) for this state to be cleared.
- alarm-only — The packet is forwarded, an alarm is generated but the source MAC is not learned on the SAP, spoke SDP or mesh SDP.
- discard-frame — The packet is discarded and an alarm generated. The frequency of alarm generation is fixed to be at most one alarm per MAC address per FP per 10 minutes in a given VPLS service. This parameter is only applicable to automatically protected MAC addresses.

When the **restrict-protected-src** is enabled on an SHG the action only applies to the associated SAPs (no action is taken by default for spoke SDPs in the SHG) and is displayed in the SAP show output as the oper state unless it is overridden by the configuration of **restrict-protected-src** on the SAP itself. In order to enable this function for spoke SDPs within a SHG, the **restrict-protected-src** must be enabled explicitly under the spoke SDP. If required, **restrict-protected-src** can also be enabled explicitly under specific SAPs within the SHG.

When this command is applied or removed, with either the **alarm-only** or **discard-frame** parameters, the MAC addresses are cleared from the related object.

The use of **restrict-protected-src discard-frame** is mutually exclusive with the configuration of manually protected MAC addresses within a given VPLS.

The **alarm-only** parameter is not supported on the 7750 SR-a, 7750 SR-1e/2e/3e.

The **no** form of the command reverts to the default.

Default no restrict-protected-src

Parameters **alarm-only** — Specifies that the packet is forwarded, an alarm is generated but the source MAC is not learned on the SAP, spoke SDP or mesh SDP,

Default no alarm-only

discard-frame — Specifies that the packet is discarded and an alarm generated. The frequency of alarm generation is fixed to be at most one alarm per MAC address per FP per 10 minutes within a given VPLS service.

Default no discard-frame

restrict-unprotected-dst

Syntax **restrict-unprotected-dst**
no restrict-unprotected-dst

Context config>service>pw-template>split-horizon-group
config>service>vpls>split-horizon-group
config>service>vpls>sap

Description This command indicates how the system will forward packets destined to an unprotected MAC address, either manually added using the **mac-protect** command or automatically added using the **auto-learn-mac-protect** command. While enabled all packets entering the configured SAP or SAPs within a split-horizon-group (but not spoke or mesh SDPs) is verified to contain a protected destination MAC address. If the packet is found to contain a non-protected destination MAC, it is discarded. Detecting a non-protected destination MAC on the SAP will not cause the SAP to be placed in the operationally down state. No alarms are generated.

If the destination MAC address is unknown, even if the packet is entering a restricted SAP, with **restrict-unprotected-dst** enabled, it is flooded.

The **no** form of the command reverts to the default.

Default no restrict-unprotected-dst

mac-pinning

Syntax [**no**] **mac-pinning**

Context config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp

Description This command disables re-learning of MAC addresses on other SAPs within the VPLS. The MAC address will remain attached to a given SAP for duration of its age-timer.

The age of the MAC address entry in the FDB is set by the age timer. If **mac-aging** is disabled on a given VPLS service, any MAC address learned on a SAP or SDP with **mac-pinning** enabled will remain in the FDB on this SAP or SDP forever.

Every event that would otherwise result in re-learning is logged (MAC address; original-SAP; new-SAP).

The **no** form of the command enables re-learning of MAC addresses.

Note: MAC addresses learned during DHCP address assignment (DHCP snooping enabled) are not impacted by this command. MAC-pinning for such addresses is implicit.

Default no mac-pinning — When a SAP or spoke SDP is part of a Residential Split Horizon Group (RSHG), MAC pinning is activated at creation of the SAP. Otherwise MAC pinning is not enabled by default.

7.3.2.1.3 ARP Handling Commands

arp-reply-agent

Syntax **arp-reply-agent** [**sub-ident**]
no arp-reply-agent

Context config>service>vpls>sap

Description This command enables a special ARP response mechanism in the system for ARP requests destined to static or dynamic hosts associated with the SAP. The system responds to each ARP request using the host's MAC address as the both the source MAC address in the Ethernet header and the target hardware address in the ARP header.

ARP replies and requests received on a SAP with **arp-reply-agent** enabled is evaluated by the system against the anti-spoof filter entries associated with the ingress SAP (if the SAP has anti-spoof filtering enabled). ARPs from unknown hosts on the SAP is discarded when anti-spoof filtering is enabled.

The ARP reply agent only responds if the ARP request enters an interface (SAP, spoke SDP or mesh SDP) associated with the VPLS instance of the SAP.

A received ARP request that is not in the ARP reply agent table is flooded to all forwarding interfaces of the VPLS capable of broadcast except the ingress interface while honoring split-horizon constraints.

Static hosts can be defined on the SAP using the **host** command. Dynamic hosts are enabled on the system by enabling the **lease-populate** command in the SAP's **dhcp** context. In the event that both a static host and a dynamic host share the same IP and MAC address, the VPLS ARP reply agent will retain the host information until both the static and dynamic information are removed. In the event that both a static and dynamic host share the same IP address, but different MAC addresses, the VPLS ARP reply agent is populated with the static host information.

The **arp-reply-agent** command will fail if an existing static host on the SAP does not have both MAC and IP addresses specified. Once the ARP reply agent is enabled, creating a static host on the SAP without both an IP address and MAC address will fail.

The ARP-reply-agent may only be enabled on SAPs supporting Ethernet encapsulation.

The **no** form of the command disables ARP-reply-agent functions for static and dynamic hosts on the SAP.

Default no arp-reply-agent

Parameters *sub-ident* — Configures the arp-reply-agent to discard ARP requests received on the SAP that are targeted for a known host on the same SAP with the same subscriber identification,

Hosts are identified by their subscriber information. For DHCP subscriber hosts, the subscriber hosts, the subscriber information is configured using the optional subscriber parameter string.

When arp-reply-agent is enabled with **sub-ident**:

- If the subscriber information for the destination host exactly matches the subscriber information for the originating host and the destination host is known on the same SAP as the source, the ARP request is silently discarded.
- If the subscriber information for the destination host or originating host is unknown or undefined, the source and destination hosts are not considered to be the same subscriber. The ARP request is forwarded outside the SAP's Split Horizon Group.
- When **sub-ident** is not configured, the arp-reply-agent does not attempt to identify the subscriber information for the destination or originating host and will not discard an ARP request based on subscriber information.

arp-populate

Syntax [**no**] **arp-populate**

Context config>service>ies>if
config>service>ies>sub-if>grp-if

Description This command, when enabled, disables dynamic learning of ARP entries. Instead, the ARP table is populated with dynamic entries from the DHCP lease state table (enabled with **lease-populate**), and optionally with static entries entered with the **host** command.

Enabling the **arp-populate** command will remove any dynamic ARP entries learned on this interface from the ARP cache.

The **arp-populate** command will fail if an existing static ARP entry exists for this interface.

The **arp-populate** command will fail if an existing static subscriber host on the SAP does not have both MAC and IP addresses specified.

Once **arp-populate** is enabled, creating a static subscriber host on the SAP without both an IP address and MAC address will fail.

When **arp-populate** is enabled, the system will not send out ARP requests for hosts that are not in the ARP cache. Only statically configured and DHCP learned hosts are reachable through an IP interface with **arp-populate** enabled. The **arp-populate** command can only be enabled on IES and VPRN interfaces supporting Ethernet encapsulation (VPRN is supported by the 7750 SR only).

The **no** form of the command disables ARP cache population functions for static and dynamic hosts on the interface. All static and dynamic host information for this interface is removed from the system's ARP cache.

Default no arp-populate

arp-timeout

Syntax **arp-timeout** *seconds*
no arp-timeout

Context config>service>ies>if
config>service>ies>sub-if>grp-if

Description This command configures the minimum time in seconds an ARP entry learned on the IP interface is stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If **arp-timeout** is set to a value of zero seconds, ARP aging is disabled.

When the **arp-populate** and **lease-populate** commands are enabled on an IES interface, the ARP table entries will no longer be dynamically learned, but instead by snooping DHCP ACK message from a DHCP server. In this case the configured **arp-timeout** value has no effect.

The default value for **arp-timeout** is 14400 seconds (4 hours).

The **no** form of this command restores **arp-timeout** to the default value.

Default arp-timeout 14400

Parameters *seconds* — Specifies the minimum number of seconds a learned ARP entry is stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 to 65535

authentication-policy

Syntax **authentication-policy** *name*
no authentication-policy

Context config>service>ies>sub-if>grp-if

Description This command assigns a RADIUS authentication policy to the interface.

The **no** form of this command removes the policy name from the group interface configuration.

Default no authentication-policy

Parameters *name* — Specifies the authentication policy name

local-proxy-arp

Syntax [**no**] **local-proxy-arp**

Context config>service>ies>if
config>service>ies>sub-if>grp-if

Description This command enables local proxy ARP. When local proxy ARP is enabled on an IP interface, the system responds to all ARP requests for IP addresses belonging to the subnet with its own MAC address, and thus comes the forwarding point for all traffic between hosts in that subnet.

When local-proxy-arp is enabled, ICMP redirects on the ports associated with the service are automatically blocked.

The **no** form of the command reverts to the default.

Default no local-proxy-arp

remote-proxy-arp

Syntax [**no**] **remote-proxy-arp**

Context config>service>ies>if
config>service>ies>sub-if>grp-if

Description This command enables remote proxy ARP on the interface.

Remote proxy ARP is similar to proxy ARP. It allows the router to answer an ARP request on an interface for a subnet that is not provisioned on that interface. This allows the router to forward to the other subnet on behalf of the requester. To distinguish remote proxy ARP from local proxy ARP, local proxy ARP performs a similar function but only when the requested IP is on the receiving interface.

The **no** form of the command reverts to the default.

Default no remote-proxy-arp

qos-route-lookup

Syntax [**no**] **qos-route-lookup**

Context config>service>ies>sub-if>grp-if
config>service>ies>sub-if>grp-if>ipv6

Description This command enables or disables QoS route lookup for the interface.

The **no** form of the command reverts to the default.

Default no qos-route-lookup

proxy-arp-policy

Syntax **proxy-arp-policy** *policy-name* [*policy-name*]
no proxy-arp-policy

Context config>service>ies>if
config>service>ies>sub-if>grp-if

Description This command specifies an existing policy-statement to analyze match and action criteria that controls the flow of routing information to and from a given protocol, set of protocols, or a particular neighbor.

The **no** form of the command removes the policy name from the configuration.

Default no proxy-arp-policy

Parameters *policy-name* — Specifies the export route policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

The specified name(s) must already be defined.

static-arp

Syntax **static-arp** *ip-address ieee-mac-address*
no static-arp *ip-address* [*ieee-mac-address*]

Context config>service>ies>if

Description This command configures a static address resolution protocol (ARP) entry associating a subscriber IP address with a MAC address for the core router instance. This static ARP appears in the core routing ARP table. A static ARP can only be configured if it exists on the network attached to the IP interface.

If an entry for a particular IP address already exists and a new MAC address is configured for the IP address, the existing MAC address is replaced with the new MAC address.

The **no** form of the command removes a static ARP entry.

Default none

Parameters *ip-address* — Specifies the IP address for the static ARP in IP address dotted decimal notation.

ieee-mac-address — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

7.3.2.2 Show Commands

Note: The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

arp

Syntax **arp** [*ip-address*] [**mac** *ieee-address*] [**sap** *sap-id*] [**interface** *ip-int-name*] [**sdp** *sdp-id:vc-id*] [**summary**]

Context show>service>id

Description This command displays the ARP cache entries for this service.

Parameters *ip-address* — Specifies the IP address in IP address dotted decimal notation.

ieee-mac-address — Specifies the 48-bit MAC address in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

sap-id — Specifies a Service Access Point (SAP) within a service.

ip-int-name — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed between double quotes.

sdp-id — Specifies the SDP identifier. Allowed values are integers in the range of 1 and 17407 for existing SDPs.

vc-id — Specifies the virtual circuit identifier.

Values 1 to 4294967295

Output The following output is an example of service ID information.

Sample Output

```
A:ALA-A# show service id 100 base
===============================================================================
ARP Table
===============================================================================
IP Address      MAC Address       Type    Age       Interface         Port
-------------------------------------------------------------------------------
10.1.0.1        00:00:66:66:66:01 Other   00h00m00s ies-10-101.1.0.1 1/1/4
239.1.1.2       00:00:5e:00:01:64 Other   00h00m00s ies-10-239.1.1.2 1/1/3
239.1.1.201     00:00:22:2e:a5:61 Static  00h00m00s ies-10-239.1.1.2 1/1/3
239.1.1.202     00:00:22:2e:a5:62 Static  00h00m00s ies-10-239.1.1.2 1/1/3
===============================================================================
A:ALA-A#
```

Table 69 **Security ARP Field Descriptions**

Label	Description
IP Address	Specifies the IP address of the ARP cache entry.
MAC Address	Specifies the MAC address associated with the IP address.
Type	Other — Learned through normal ARP queries. Static — Configured by **static-arp** commands. Managed — Learned from DHCP snooping or configured by **host** commands.
Age	Indicates age of the ARP entry.
Interface	Indicates the name of the IP interface.
Port	Indicates the port upon which the entry was learned.

7.3.2.3 Tools Commands

authentication-server-check

Syntax **authentication-server-check server-address** *ip-address* [**port** *port*] {{**user-name** *user-name* **password** *password*} | **attr-from-file** *file-url*} **secret** *key* [**source-address** *ip-address*] [**timeout** *seconds*] [**router** *router-instance* | **service-name** *service-name*] [**debug**]

Context tools>perform>security

Description This command is used to perform an authentication check on the RADIUS server.

Parameters **server-address** *ip-address* — Specifies the IP address of the RADIUS server.

port — Specifies the port of the RADIUS server.

Values 1 to 65535

user-name — Specifies the user name, up to 253 characters.

password — Specifies the user password, up to 64 characters.

file-url — specifies to attach additional VSAs from a file.

Values

local-url:	[*clfash-id*] [*file-path*]
	200 characters maximum, including cflash-id directory length up to 99 characters each
remote-url:	[{ftp://|tftp://}<login>:<pswd>@<remote-locn>/][<file-path>]
	255 characters maximum, directory length maximum 99 characters each
remote-locn:	[<hostname> | <ipv4-address> | <ipv6-address>]
ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x[-interface]
	x:x:x:x:x:x:d.d.d.d[-interface]
	x - [0..FFFF]H
	d - [0..255]D
	interface - 32 chars max, for link local addresses
cflash-id:	cf1:, cf1-A:, cf1-B:, cf2:, cf2-A:, cf2-B:, cf3:, cf3-A:, cf3-B:

key — Specifies the secret key shared between the RADIUS client and server, up to 64characters.

source address *ip-address* — Specifies the source address used for the access-request message.

Values

ipv4-address	a.b.c.d
ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D

seconds — Specifies the time to wait for the RADIUS server response.

Values 1 to 90

router-instance — Specifies the router instance that should be used by the access request message.

Values
router-name | *vprn-svc-id*

router-name	Base, management
	Default - Base
vprn-svc-id	1 to 2147483647

service-name — Specifies the service name that should be used by the access request message, up to 64 characters.

debug — Specifies to display the full debug message, including all VSAs and full hex sting, of the RADIUS reply message.

8 Triple Play Multicast

8.1 Introduction to Multicast

IP multicast provides an effective method of many-to-many communication. Delivering unicast datagrams is fairly simple. Normally, IP packets are sent from a single source to a single recipient. The source inserts the address of the target host in the IP header destination field of an IP datagram, intermediate routers (if present) simply forward the datagram towards the target in accordance with their respective routing tables.

Sometimes distribution needs individual IP packets be delivered to multiple destinations (like audio or video streaming broadcasts). Multicast is a method of distributing datagrams sourced from one (or possibly more) host(s) to a set of receivers that may be distributed over different (sub) networks. This makes delivery of multicast datagrams significantly more complex.

Multicast sources can send a single copy of data using a single address for the entire group of recipients. The routers between the source and recipients route the data using the group address route. Multicast packets are delivered to a multicast group. A multicast group specifies a set of recipients who are interested in a particular data stream and is represented by an IP address from a specified range. Data addressed to the IP address is forwarded to the members of the group. A source host sends data to a multicast group by specifying the multicast group address in datagram's destination IP address. A source does not have to register in order to send data to a group nor do they need to be a member of the group.

Routers and Layer 3 switches use the Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) to manage membership for a multicast session. When a host wants to receive one or more multicast sessions, it sends a join message for each multicast group it wants to join. When a host wants to leave a multicast group, it sends a leave message.

8.2 Multicast in the Broadband Service Router

This section describes the multicast protocols employed when a Nokia router is used as a Broadband Service Router (BSR) in a Triple Play aggregation network.

The protocols used are:

- Internet Group Management Protocol (Internet Group Management Protocol)
- Multicast Listener Discovery (Multicast Listener Discovery)
- Source Specific Multicast Groups (Internet Group Management Protocol)
- Protocol Independent Multicast (Sparse Mode) (Protocol Independent Multicast Sparse Mode (PIM-SM))

8.2.1 Internet Group Management Protocol

Internet Group Management Protocol (IGMP) is used by IPv4 hosts and routers to report their IP multicast group memberships to neighboring multicast routers. A multicast router keeps a list of multicast group memberships for each attached network, and a timer for each membership.

Multicast group memberships include at least one member of a multicast group on a given attached network, not a list of all of the members. With respect to each of its attached networks, a multicast router can assume one of two roles, querier or non-querier. There is normally only one querier per physical network.

A querier issues two types of queries, a general query and a group-specific query. General queries are issued to solicit membership information with regard to any multicast group. Group-specific queries are issued when a router receives a leave message from the node it perceives as the last group member remaining on that network segment.

Hosts wanting to receive a multicast session issue a multicast group membership report. These reports must be sent to all multicast enabled routers.

8.2.1.1 IGMP Versions and Interoperability Requirements

If routers run different versions of IGMP, they negotiates the lowest common version of IGMP that is supported on their subnet and operate in that version.

Version 1 — Specified in RFC-1112, *Host extensions for IP Multicasting*, was the first widely deployed version and the first version to become an Internet standard.

Version 2 — Specified in RFC-2236, *Internet Group Management Protocol*, added support for low leave latency, that is, a reduction in the time it takes for a multicast router to learn that there are no longer any members of a particular group present on an attached network.

Version 3 —Specified in RFC-3376, *Internet Group Management Protocol*, adds support for source filtering, that is, the ability for a system to report interest in receiving packets only from specific source addresses, as required to support Source-Specific Multicast (See Source Specific Multicast (SSM)), or from all but specific source addresses, sent to a particular multicast address.

IGMPv3 must keep state per group per attached network. This group state consists of a filter-mode, a list of sources, and various timers. For each attached network running IGMP, a multicast router records the desired reception state for that network.

8.2.1.2 IGMP Version Transition

Nokia's SRs are capable of interoperating with routers and hosts running IGMPv1, IGMPv2, and/or IGMPv3. *Draft-ietf-magma-igmpv3-and-routing-0x.txt* explores some of the interoperability issues and how they affect the various routing protocols.

IGMP version 3 specifies that if at any point a router receives an older version query message on an interface that it must immediately switch into a compatibility mode with that earlier version. Since none of the previous versions of IGMP are source aware, should this occur and the interface switch to Version 1 or 2 compatibility mode, any previously learned group memberships with specific sources (learned from the IGMPv3 specific INCLUDE or EXCLUDE mechanisms) must be converted to non-source specific group memberships. The routing protocol then treats this as if there is no EXCLUDE definition present.

8.2.2 Multicast Listener Discovery

Multicast Listener Discovery (MLD) is used by IPv6 hosts and routers to report their IP multicast group memberships to neighboring multicast routers. A multicast router keeps a list of multicast group memberships for each attached network, and a timer for each membership. Multicast group memberships include at least one member of a multicast group on a given attached network, not a list of all of the members. With respect to each of its attached networks, a multicast router can assume one of two roles, querier or non-querier. There is normally only one querier per physical network.

A querier issues two types of queries, a general query and a group-specific query. General queries are issued to solicit membership information with regard to any multicast group. Group-specific queries are issued when a router receives a leave message from the node it perceives as the last group member remaining on that network segment.

Hosts wanting to receive a multicast session issue a multicast group membership report. These reports must be sent to all multicast enabled routers.

8.2.2.1 MLD Versions and Interoperability Requirements

If routers run different versions of MLD, they negotiates the lowest common version of MLD that is supported on their subnet and operate in that version.

Version 1 — Specified in RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*, was the first deployed version and included low leave latency, that is, a reduction in the time it takes for a multicast router to learn that there are no longer any members of a particular group present on an attached network.

Version 2 — Specified in RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*, adds support for source filtering, that is, the ability for a system to report interest in receiving packets only from specific source addresses, as required to support Source-Specific Multicast.

Multicast (SSM)), or from all but specific source addresses, sent to a particular multicast address. MLDv2 must keep state per group per attached network. This group state consists of a filter mode, a list of sources, and various timers. For each attached network running MLD, a multicast router records the desired reception state for that network.

8.2.2.2 Source Specific Multicast Groups

IGMPv3and MLDv2 permits a receiver to join a group and specify that it only wants to receive traffic for a group if that traffic comes from a particular source. If a receiver does this, and no other receiver on the LAN requires all the traffic for the group, then the Designated Router (DR) can omit performing a (*,G) join to set up the shared tree, and instead issue a source-specific (S,G) join only.

For IPv4, the range of multicast addresses from 232.0.0.0 to 232.255.255.255 is currently set aside for source-specific multicast. For groups in this range, receivers should only issue source specific IGMPv3 joins. If a PIM router receives a non-source-specific join for a group in this range, it should ignore it.

For IPv6, the multicast prefix FF3x::/32 is currently set aside for source-specific multicast. For groups in this range, receivers should only issue source specific MLDv2 joins. If a PIM router receives a non-source-specific join for a group in this range, it should ignore it.

A Nokia PIM router must silently ignore a received (*, G) PIM join message where G is a multicast group address from the multicast address group range that has been explicitly configured for SSM. This occurrence should generate an event. If configured, the IGMPv2 (MLDv1for IPv6) request can be translated into IGMPv3 (MLDv2 for IPv6). The SR allows for the conversion of an IGMPv2 (*,G) request into a IGMPv3 (S,G) request based on manual entries. A maximum of 32 SSM ranges is supported.

IGMPv3 and MLDv2 also permits a receiver to join a group and specify that it only wants to receive traffic for a group if that traffic does not come from a specific source or sources. In this case, the DR performs a (*,G) join as normal, but can combine this with a prune for each of the sources the receiver does not wish to receive.

8.2.3 Protocol Independent Multicast Sparse Mode (PIM-SM)

PIM-SM leverages the unicast routing protocols that are used to create the unicast routing table: OSPF, IS-IS, BGP, and static routes. Because PIM uses this unicast routing information to perform the multicast forwarding function it is effectively IP protocol independent. Unlike DVMRP, PIM does not send multicast routing tables updates to its neighbors.

PIM-SM uses the unicast routing table to perform the Reverse Path Forwarding (RPF) check function instead of building up a completely independent multicast routing table.

PIM-SM only forwards data to network segments with active receivers that have explicitly requested the multicast group. PIM-SM in the ASM model initially uses a shared tree to distribute information about active sources. Depending on the configuration options, the traffic can remain on the shared tree or switch over to an optimized source distribution tree. As multicast traffic starts to flow down the shared tree, routers along the path determine if there is a better path to the source. If a more direct path exists, then the router closest to the receiver sends a join message toward the source and then reroutes the traffic along this path.

As stated above, PIM-SM relies on an underlying topology-gathering protocol to populate a routing table with routes. This routing table is called the Multicast Routing Information Base (MRIB). The routes in this table can be taken directly from the unicast routing table, or it can be different and provided by a separate routing protocol such as MBGP. Regardless of how it is created, the primary role of the MRIB in the PIM-SM protocol is to provide the next hop router along a multicast-capable path to each destination subnet. The MRIB is used to determine the next hop neighbor to whom any PIM join/prune message is sent. Data flows along the reverse path of the join messages. Thus, in contrast to the unicast RIB that specifies the next hop that a data packet would take to get to some subnet, the MRIB gives reverse-path information, and indicates the path that a multicast data packet would take from its origin subnet to the router that has the MRIB.

8.2.4 Ingress Multicast Path Management (IMPM) Enhancements

Refer to the *Advanced Configuration Guide* for more information about IMPM as well as detailed configuration examples.

IMPM allows the system to manage Layer 2 and Layer 3 IP multicast flows by sorting them into the available multicast paths through the switch fabric. The ingress multicast manager tracks the amount of available multicast bandwidth per path and the amount of bandwidth used per IP multicast stream. The following traffic is managed by IMPM when enabled:

- IPv4 and IPv6 routed multicast traffic
- VPLS IGMP snooping traffic
- VPLS PIM snooping for IPv4 traffic
- VPLS PIM snooping for IPv6 traffic (only when sg-based forwarding is configured)

Two policies define how each path should be managed, the bandwidth policy, and how multicast channels compete for the available bandwidth, the multicast information policy.

Chassis multicast planes should not be confused with IOM/IMM multicast paths. The IOM/IMM uses multicast paths to reach multicast planes on the switch fabric. An IOM/IMM may have less or more multicast paths than the number of multicast planes available in the chassis.

Each IOM/IMM multicast path is either a primary or secondary path type. The path type indicates the multicast scheduling priority within the switch fabric. Multicast flows sent on primary paths are scheduled at multicast high priority while secondary paths are associated with multicast low priority.

The system determines the number of primary and secondary paths from each IOM/IMM forwarding plane and distributes them as equally as possible between the available switch fabric multicast planes. Each multicast plane may terminate multiple paths of both the primary and secondary types.

The system ingress multicast management module evaluates the ingress multicast flows from each ingress forwarding plane and determines the best multicast path for the flow. A particular path may be used until the terminating multicast plane is “maxed” out (based on the rate limit defined in the **per-mcast-plane-capacity** commands) at which time either flows are moved to other paths or potentially blackholed (flows with the lowest preference are dropped first). In this way, the system makes the best use of the available multicast capacity without congesting individual multicast planes.

The switch fabric is simultaneously handling both unicast and multicast flows. The switch fabric uses a weighted scheduling scheme between multicast high, unicast high, multicast low and unicast low when deciding which cell to forward to the egress forwarding plane next. The weighted mechanism allows some amount of unicast and lower priority multicast (secondary) to drain on the egress switch fabric links used by each multicast plane. The amount is variable based on the number of switch fabric planes available on the amount of traffic attempting to use the fabric planes. The per-mcast-plane-capacity commands allows the amount of managed multicast traffic to be tuned to compensate for the expected available egress multicast bandwidth per multicast plane. In conditions where it is highly desirable to prevent multicast plane congestion, the per-mcast-plane-capacity commands should be used to compensate for the non-multicast or secondary multicast switch fabric traffic.

8.3 Multicast in the BSA

IP Multicast is normally not a function of the Broadband Service Aggregator (BSA) in a Triple Play aggregation network being a Layer 2 device. However, the BSA does use IGMP snooping to optimize bandwidth utilization.

8.3.1 IGMP Snooping

For most Layer 2 switches, multicast traffic is treated like an unknown MAC address or broadcast frame, which causes the incoming frame to be flooded out (broadcast) on every port within a VLAN. While this is acceptable behavior for unknowns and broadcasts, as IP Multicast hosts may join and be interested in only specific multicast groups, all this flooded traffic results in wasted bandwidth on network segments and end stations.

IGMP snooping entails using information in layer 3 protocol headers of multicast control messages to determine the processing at layer 2. By doing so, an IGMP snooping switch provides the benefit of conserving bandwidth on those segments of the network where no node has expressed interest in receiving packets addressed to the group address.

On the 7450 ESS and 7750 SR, IGMP snooping can be enabled in the context of VPLS services. The IGMP snooping feature allows for optimization of the multicast data flow for a group within a service to only those Service Access Points (SAPs) and Service Distribution Points (SDPs) that are members of the group. In fact, the 7450 ESS and 7750 SR implementation performs more than pure snooping of IGMP data, since it also summarizes upstream IGMP reports and responds to downstream queries.

The 7450 ESS and 7750 SR maintain a number of multicast databases:

- A port database on each SAP and SDP lists the multicast groups that are active on this SAP or SDP.
- All port databases are compiled into a central proxy database. Towards the multicast routers, summarized group membership reports are sent based on the information in the proxy database.
- The information in the different port databases is also used to compile the multicast forwarding information base (MFIB). This contains the active SAPs and SDPs for every combination of source router and group address (S,G), and is used for the actual multicast replication and forwarding.

When the router receives a join report from a host for a particular multicast group, it adds the group to the port database and (if it is a new group) to the proxy database. It also adds the SAP or SDP to existing (S,G) in the MFIB, or builds a new MFIB entry.

When the router receives a leave report from a host, it first checks if other devices on the SAP or SDP still want to receive the group (unless fast leave is enabled). Then it removes the group from the port database, and from the proxy database if it was the only receiver of the group. The router also deletes entries if it does not receive periodic membership confirmations from the hosts.

The fast leave feature finds its use in multicast TV delivery systems, for example. Fast Leave speeds up the membership leave process by terminating the multicast session immediately, rather than the standard procedure of issuing a group specific query to check if other group members are present on the SAP or SDP.

8.3.1.1 IGMP/MLD Message Processing

Figure 50 illustrates the basic IGMP message processing in several situations.

Figure 50 **IGMP/MLD Message Processing**

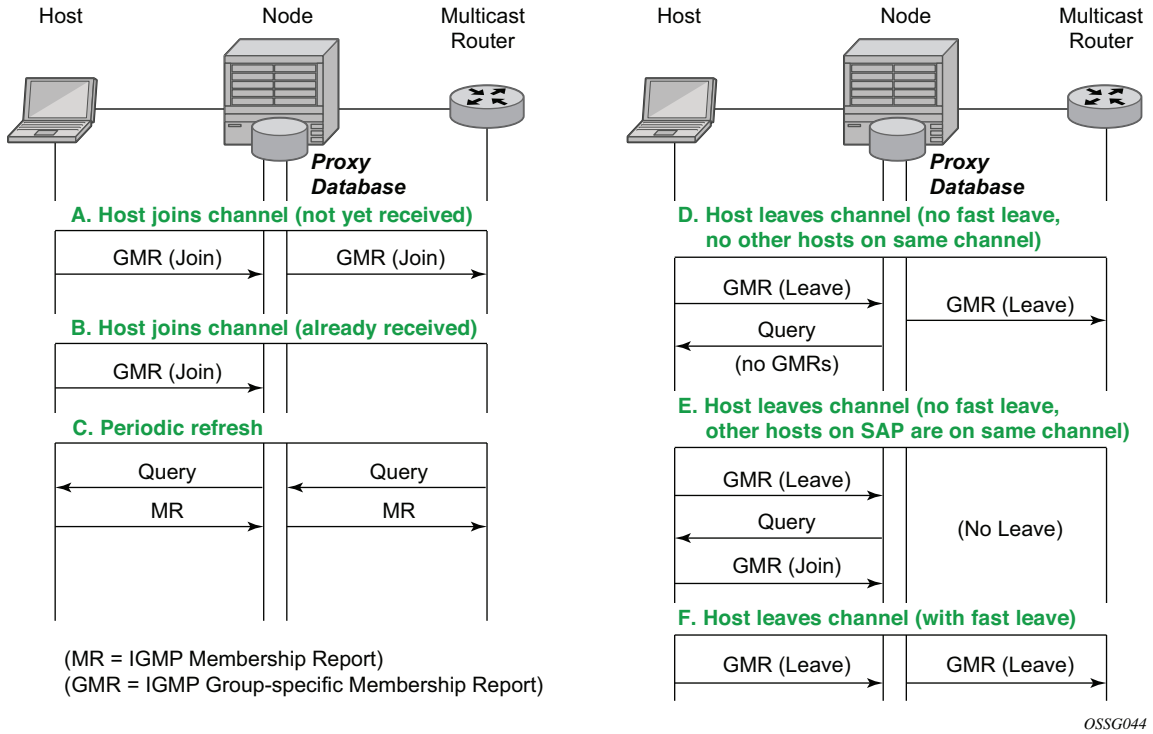

8.3.1.2 IGMP Message Processing

Scenario A: A host joins a multicast group (TV channel) which is not yet being received by other hosts on the router, and thus is not yet present in the proxy database. The 7450 ESS or 7750 SR adds the group to the proxy database and sends a new IGMP Join group-specific membership report upstream to the multicast router.

Scenario B: A host joins a channel which is already being received by one or more hosts on the router, and thus is already present in the proxy database. No upstream IGMP report is generated by the router.

Scenario C: The multicast router periodically sends IGMP queries to the router, requesting it to respond with generic membership reports. Upon receiving such a query, the router compiles a report from its proxy database and send it back to the multicast router.

In addition, the router floods the received IGMP query to all hosts (on SAPs and spoke SDPs), and updates its proxy database based on the membership reports received back.

Scenario D: A host leaves a channel by sending an IGMP leave message. If fast-leave is not enabled, the router first checks whether there are other hosts on the same SAP or spoke SDP by sending a query. If no other host responds, the 7450 ESS or 7750 SR removes the channel from the SAP. In addition, if there are no other SAPs or spoke SDPs with hosts subscribing to the same channel, the channel is removed from the proxy database and an IGMP leave report is sent to the upstream Multicast Router.

Scenario E: A host leaves a channel by sending an IGMP leave message. If fast-leave is not enabled, the router checks whether there are other hosts on the same SAP or spoke SDP by sending a query. Another device on the same SAP or spoke SDP still wishes to receive the channel and responds with a membership report. Thus the router does not remove the channel from the SAP.

Scenario F: A host leaves a channel by sending an IGMP leave report. Fast-leave is enabled, so the router will not check whether there are other hosts on the same SAP or spoke SDP but immediately removes the group from the SAP. In addition, if there are no other SAPs or spoke SDPs with hosts subscribing to the same group, the group is removed from the proxy database and an IGMP leave report is sent to the upstream multicast router.

8.3.1.3 MLD Message Processing

MLD message processing differs from IGMP. An IPv6 host can have two WAN IPv6 addresses and an IPv6 prefix. MLD messages source address are link local addresses. This makes it difficult to know if the originating host is a WAN host or a PD host. By default, all requested IPv6 (s,g) are first associated with a WAN host. If this particular WAN host disconnects or ends its IPv6 session, the (s,g) is then associated with the remaining WAN host. If there are no more WAN hosts, the (s,g) is then associated with the remaining PD host. The (s,g) is always transferred to the remaining IPv6 host until there are no more report replies to corresponding to the queries. Scenarios A — F will not differ for IPv6 hosts.

8.3.1.4 IGMP/MLD Filtering

A provider may want to block receive or transmit permission to individual hosts or a range of hosts. To this end, the 7450 ESS and 7750 SR support IGMP/MLD filtering. Two types of filter can be defined:

- Filter IGMP/MLD membership reports from a particular host or range of hosts. This is performed by importing an appropriately defined routing policy into the SAP or spoke SDP.
- Filter to prevent a host from transmitting multicast streams into the network. The operator can define a data-plane filter (ACL) which drops all multicast traffic, and apply this filter to a SAP or spoke SDP.

8.3.2 Multicast VPLS Registration (MVR)

Multicast VPLS Registration (MVR) is a bandwidth optimization method for multicast in a broadband services network. MVR allows a subscriber on a port to subscribe and unsubscribe to a multicast stream on one or more network-wide multicast VPLS instances.

MVR assumes that subscribers join and leave multicast streams by sending IGMP join and leave messages. The IGMP leave and join message are sent inside the VPLS to which the subscriber port is assigned. The multicast VPLS is shared in the network while the subscribers remain in separate VPLS services. Using MVR, users on different VPLS cannot exchange any information between them, but still multicast services are provided.

On the MVR VPLS, IGMP snooping must be enabled. On the user VPLS, IGMP snooping and MVR work independently. If IGMP snooping and MVR are both enabled, MVR reacts only to join and leave messages from multicast groups configured under MVR. Join and leave messages from all other multicast groups are managed by IGMP snooping in the local VPLS. This way, potentially several MVR VPLS instances could be configured, each with its own set of multicast channels.

MVR by proxy — In some situations, the multicast traffic should not be copied from the MVR VPLS to the SAP on which the IGMP message was received (standard MVR behavior) but to another SAP. This is called MVR by proxy.

Figure 51 shows a MVR and MVR by proxy configuration.

Figure 51 **MVR and MVR by Proxy**

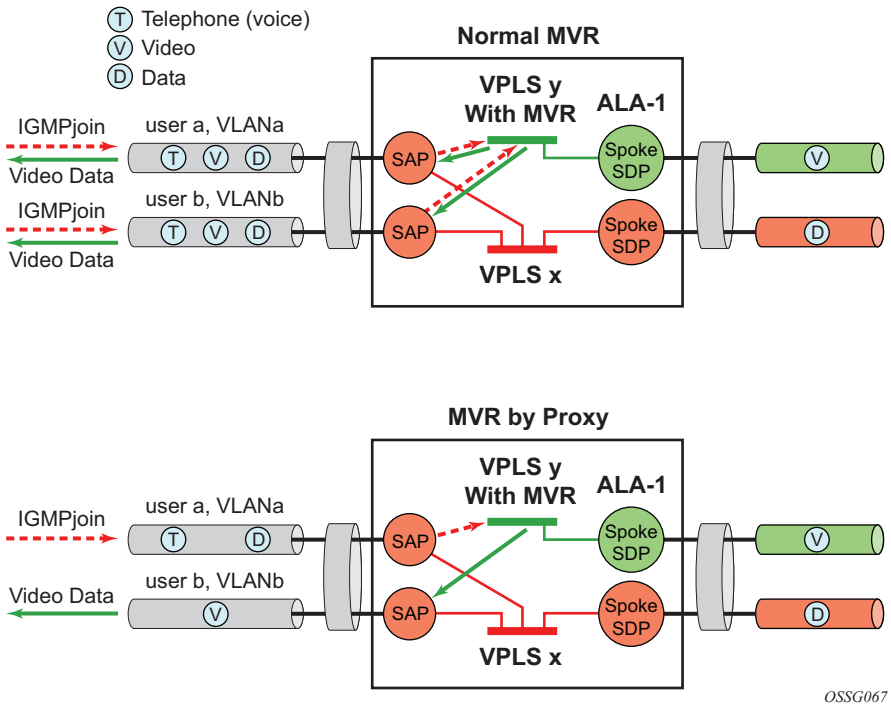

8.3.3 Layer 3 Multicast Load Balancing

Layer 3 multicast load balancing establishes a more efficient distribution of Layer 3 multicast data over ECMP and LAG links. Operators have the option to redistribute multicast groups over ECMP and/or LAG links if the number of links changes either up or down.

When implementing this feature, there are several considerations. When multicast load balancing is not configured, the distribution remains as is. Multicast load balancing is based on the number of “s,g” groups. This means that bandwidth considerations are not taken into account. The multicast groups are distributed over the available links as joins are processed. When link failure occurs, the load is distributed on the failed channel to the remaining channels so multicast groups are evenly distributed over the remaining links. When a link is added (or failed link returned) all multicast joins on the added link(s) are allocated until a balance is achieved.

When multicast load balancing is configured, but the channels are not found in the multicast-info-policy, then multicast load balancing is based on the number of “s,g” groups. This means that bandwidth considerations are not taken into account. The multicast groups are distributed over the available links as joins are processed. The multicast groups are evenly distributed over the remaining links. When link failure occurs, the load is distributed on the failed channel to the remaining channels. When a link is added (or failed link returned) all multicast joins on the added link(s) are allocated until a balance is achieved. A manual redistribute command enables the operator to re-evaluate the current balance and, if required, move channels to different links to achieve a balance. A timed redistribute parameter allows the system to automatically, at regular intervals, redistribute multicast groups over available links. If no links have been added or removed from the ECMP/LAG interface, then no redistribution is attempted.

When multicast load balancing is configured, multicast groups are distributed over the available links as joins are processed based on bandwidth configured for the specified group address. If the bandwidth is not configured for the multicast stream then the configured default value is used.

If link failure occurs, the load is distributed on the failed channel to the remaining channels. The bandwidth required over each individual link is evenly distributed over the remaining links.

When an additional link is available for a given multicast stream, then it is considered in all multicast stream additions applied to the interface. This multicast stream is included in the next scheduled automatic rebalance run. A rebalance run re-evaluates the current balance with regard to the bandwidth utilization and if required, move multicast streams to different links to achieve a balance.

A rebalance, either timed or executing the **mc-ecmp-rebalance** command, should be administered gradually in order to minimize the effect of the rebalancing process on the different multicast streams. If multicast re-balancing is disabled and subsequently enabled, keeping with the rebalance process, the gradual and least invasive method is used to minimize the effect of the changes to the customer.

By default multicast load balancing over ECMP links is enabled and set at 30 minutes.

The rebalance process can be executed as a low priority background task while control of the console is returned to the operator. When multicast load rebalancing is not enabled, then ECMP changes will not be optimized, however, when a link is added occurs an attempt is made to balance the number of multicast streams on the available ECMP links. This however may not result in balanced utilization of ECMP links.

Only a single **mc-ecmp-rebalance** command can be executed at any given time, if a rebalance is in progress and the command is entered, it is rejected with the message saying that a rebalance is already in progress. A low priority event is generated when an actual change for a given multicast stream occurs as a result of the rebalance process.

8.3.4 IGMP State Reporter

The target application for this feature is linear TV delivery. In some countries, wholesale Service Providers are obligated by the government regulation to provide information about channel viewership per subscriber to retailers.

A service provider (wholesaler or retailer) my use this information for:

- billing purposes
- market research/data mining to gain view into the most frequently watched channels, duration of the channel viewing, frequency of channel zapping by the time of the day, and so on

The information about channel viewership is based on IGMP states maintained per each subscriber host. Each event related to the IGMP state creation is recorded and formatted by the IGMP process. The formatted event is then sent to another task in the system (Exporter), which allocates a TX buffer and start a timer.

The event is then be written by the Exporter into the buffer. The buffer in essence corresponds to the packet that contains a single event or a set of events. Those events are transported as data records over UDP transport to an external collector node. The packet itself has a header followed by a set of TLV type data structures, each describing a unique filed within the IGMP event.

The packet is transmitted when it reaches a preconfigured size (1400 bytes), or when the timer expires, whichever comes first.

Note: The timer started when the buffer was initially created.

The receiving end (collector node) accepts the data on the destination UDP port. It must be aware of the data format so that it can interpret incoming data accordingly. The implementation details of the receiving node are outside of the scope of this description and are left to the network operator.

The IGMP state recording per subscriber host must be supported for hosts which are replicating multicast traffic directly as well as for those host that are only keeping track of IGMP states for the HQoS Adjustment purpose. The latter is implemented via redirection and not the Host Tracking (HT) feature as originally proposed. The IGMP reporting must differentiate events between direct replication and redirection.

It further distinguish events that are related to denial of IGMP state creation (due to filters, MCAC failure, and so on) and the ones that are related to removal of an already existing IGMP state in the system.

8.3.4.1 IGMP Data Records

Each IGMP state change generates a data record that is formatted by the IGMP task and written into the buffer. IGMP state transitions configured statically through CLI are not reported.

In order to minimize the size of the records when transported over the network, most fields in the data record are HEX coded (as opposed to ASCII descriptive strings).

Each data record has a common header as shown in Figure 52:

Figure 52 **Common IGMP Data Record Header**

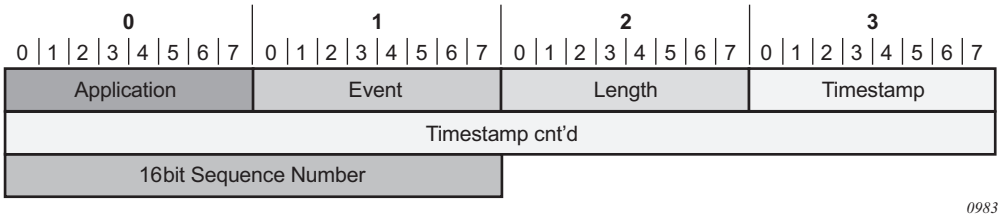

Application:

- 0x01 - IGMP

- 0x02 - IGMP Host Tracking Event:

Event:

- Related to denial of a new state creation:
  - 0x01 – Join
  - 0x02 – (Join_Deny_Filter) Join denied due to filtering via an import policy
  - 0x03 – (Join_Deny_CAC) Join denied due to MCAC
  - 0x04 – (Join_Deny_MaxGrps) Join denied due to maximum groups per host limit reached
  - 0x05 – (Join_Deny_MaxSrcs) Join denied due to maximum sources limit reached
  - 0x06 - Join (Join_Deny_SysErr) Join denied due to an internal error (for example: out of memory)
  - Related to removal of an existing IGMP state:
  - 0x07 - (Drop_Leave_Rx) IGMP state is removed due to the Leave message
  - 0x08 - (Drop_Expiry) IGMP state is removed due to time out (by default 2*query_interval + query_response_interval = 260sec)
  - 0x09 - (Drop_Filter) IGMP state is removed due to filter (import policy) change
  - 0x0A - (Drop_CfgChange) IGMP state is removed due to configuration change (clear grp, intf shutdown, PPPoE session goes unexpectedly down)
  - 0x0B - (Drop_CAC) an existing stream is stopped due to configuration change in MCAC

Length:

- The length of the entire data record (including the header and TLVs) in octets.

16 bit Sequence Number

- Since IGMP Reporting is based on connectionless transport (UDP), a 16 bit sequence numbers are used in each data record so that data loss in the network can be tracked.
- The 16 bit sequence number is located after the time stamp field. The sequence numbers increases sequentially from 0 — 65535 and then rollover back to 0.

Timestamp:

- Timestamp is in Unix format (32 bit integer in seconds since 01/01/1970) plus an extra 8 bits for 10msec resolution.

TLVs describing the IGMP state record has the following structure shown in Figure 53 and the data record field descriptions in Table 70:

Figure 53 **Data Record Field TLV Structure**

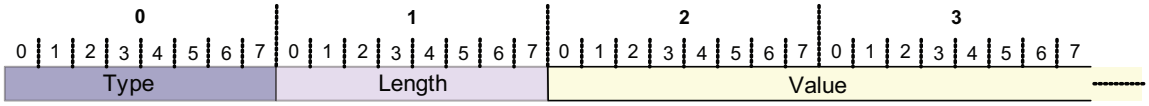

Table 70 **Data Record Field Description**

Type	Description	Encoding/Length	Mandatory/Optional
0x02	Subscriber ID	ASCII	M
0x03	Sub Host IP	4 Bytes IPv4	M
0x04	Mcast Group IP	4 Bytes IPv4	M
0x05	Mcast Source IP	4 Bytes IPv4	M
0x06	Host MAC	6 Bytes	M
0x07	PPPoE Session-ID	2 Bytes	M
0x08	Service ID	4 Bytes	M
0x09	SAP ID	ASCII	M
0x0A	Redirection vRtrId	4 Bytes	M
0x0B	Redirection ifIndex	4 Bytes	M

The *redirection destination* TLV is a mandatory TLV that is sent only in cases where redirection is enabled. It contains two 32 bit integer numbers. The first number identifies the VRF where IGMPs are redirected; the second number identifies the interface index.

Optional fields can be included in the data records according to the configuration.

In IGMPv3, if an IGMP message (Join or Leave) contains multiple multicast groups or a multicast group contains multiple IP sources, only a single event is generated per group-source combination. In other words, data records are transmitted with a single source IP address and multiple mcast group addresses or a single multicast group address with multiple source IP addresses, depending on the content of the IGMP message. (*,G)

8.3.4.2 Transport Mechanism

Data is transported via UDP socket. Destination IP address, the destination port and the source IP address are configurable. The default UDP source and destination port number is 1037.

Upon the arrival of an IGMP event, the Exporter allocates a buffer for the packet (if not already allocated) and starts writing the events into the buffer (packet). Along with the initial buffer creation, a timer is started. The trigger for the transmission of the packet is either the TX buffer being filled up to 1400B (hard coded value), or the timer expiry, whichever comes first.

The source IP address is configurable within GRT (by default system IP), and the destination IP address can be reachable only via GRT. The source IP address is modified via **system>security>source-address>application** CLI hierarchy.

The receiving end (the collector node) collects the data and process them according to the formatting rules defined in this document. The capturing and processing of the data on the collector node is outside of the context of this description.

The processing node should have sufficient resources to accept and process packets that contain information about every IGMP state change for every host from a set of network BRASes that are transporting data to this particular collector node.

Multicast Reporter traffic is marked as BE (all 6 DSCP bits are set to 0) exiting our system.

8.3.4.3 HA Compliance

IGMP Events are synchronized between two CPMs before they are transported out of the system.

8.3.4.4 QoS Awareness

IGMP Reporter is a client of sgt-qos so that DSCP/dot1p bits can be appropriately be marked when egressing the system.

8.3.4.5 IGMP Reporting Caveats

The following are not supported:

- Regular (non-subscriber) interfaces
- SAM support as the collector device

8.4 Multicast Support over Subscriber Interfaces in Routed CO Model

Applications for multicast over subscriber interfaces in Routed CO ESM model can be divided in two main categories:

Residential customers where the driver applications are:

- IPTV in an environment with legacy non-multicasting DSLAMs
- Internet multicast where users connect to a multicast stream sourced from the Internet.

For the business customers, the main drivers are enterprise multicast and Internet multicast applications.

On multicast-capable ANs, a single copy of each multicast stream is delivered over a separate regular IP interface. AN would then perform the replication. This is how multicast would be deployed in Routed CO environment with the 7750 SR and 7450 ESS.

On legacy, non-multicast ANs, or in environments with low volume multicast traffic where it is not worth setting up a separate multicast topology (from BNG to AN), multicast replication is performed via subscriber-interfaces in the 7750 SR and 7450 ESS. There are differences in replicating multicast traffic on IPoE vs PPPoX which are described in subsequent sessions.

An example of a business connectivity model is shown in Figure 54.

Figure 54 **A Typical Business Connectivity Model**

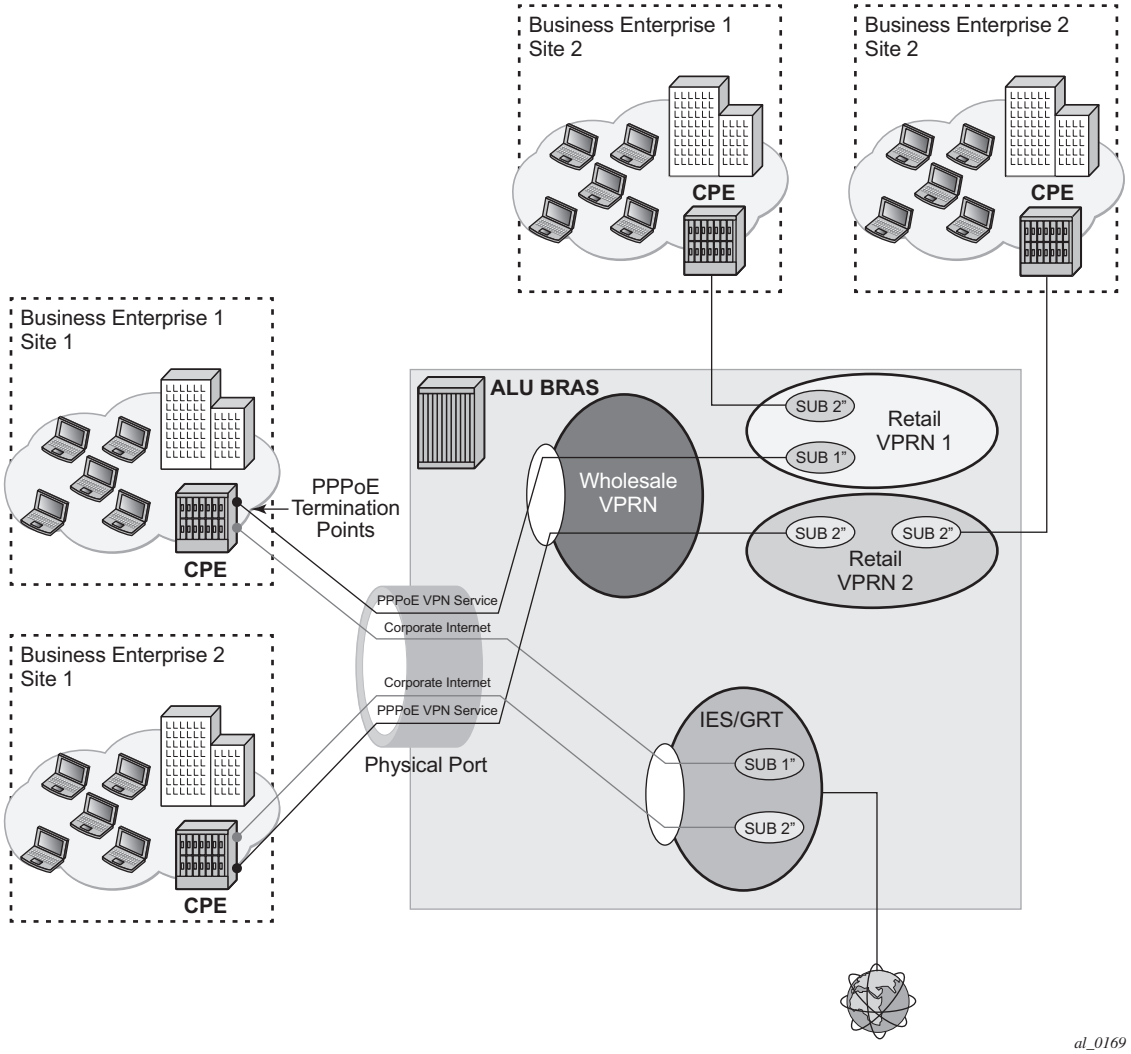

In this example, HSI is terminated in a Global Routing Table (GRT) whereas VPRN services are terminated in Wholesale/Retail VPRN fashion, with each customer using a separate VPRN.

The actual connectivity model that is deployed depends on many operational aspects that are present in the customer environment.

Multicast over subscriber-interfaces in a Routed CO model is supported for both types of hosts, IPoE and PPPoE which can be simultaneously enabled on a shared SAP.

There are some fundamental differences in multicast behavior between two host types (IPoE and PPPoX). The differences are discussed further in the next sections.

8.4.1 Multicast Over IPoE

There are several deployment scenarios for delivering multicast directly over subscriber hosts:

- 1:1 model (subscriber per VLAN/SAP) with the Access Node (AN) that is not IGMP/MLD aware.
- N:1 model (service per VLAN/SAP) with the AN in the Snooping mode.
- N:1 model with the AN in the Proxy mode.
- N:1 model with the AN that is not IGMP/MLD aware.

There are two modes of operation for subscriber multicast that can be chosen to address the above mentioned deployment scenarios:

1. Per SAP replication — A single multicast stream per group is forwarded on any given SAP. Even if the SAP has a multicast group (channel) that is registered to multiple hosts, only a single copy of the multicast stream is forwarded over this SAP. The multicast stream has a multicast destination MAC address (as opposed to unicast). IGMP/MLD states are maintained per host. This is the default mode of operation.
2. Per subscriber host replication in this mode of operation, multicast is replicated per subscriber host even if this means that multiple copies of the same stream that are forwarded over the same SAP. For example, if two hosts on the same SAP are registered to receive the same multicast group (channel), then this multicast channel is replicated twice on the same SAP. The streams have a unique unicast destination MAC address (otherwise it would not make sense to replicate the streams twice).

In all deployment scenarios and modes of operation the IGMP/MLD states per source IP address of the incoming IGMP/MLD message is maintained. This source IP address might represent a subscriber hosts or the AN (proxy mode).

For MLD, the source IP address is the host link local address. Therefore, the MLD message is associated with all IP address/prefix of the host.

8.4.1.1 Per SAP Replication Mode

In the per SAP replication mode a single copy of the multicast channel is forwarded per SAP. In other words, if a subscriber (in 1:1 mode) or a group of subscribers (in N:1 mode) have multiple hosts and all of them are subscribed to the same multicast group (watching the same channel), then only a single copy of the multicast stream for that group is sent. The destination MAC address is always a multicast MAC (there are no conversion to unicast MAC address).

IGMP/MLD states are maintained per subscriber host and per SAP.

8.4.1.1.1 Per SAP Queue

Multicast traffic over subscribers in a per-SAP replication mode flows via a SAP queue which is outside of the subscriber queues context. Sending the multicast traffic over the default SAP queue is characterized by:

- the inability to classify multicast traffic into separate subscriber queues and therefore include it natively in the subscriber HQoS hierarchy; however, multicast traffic can be classified into specific (M-)SAP queues, assuming that such queues are enabled via (M-)SAP based QoS policy
- redirection of multicast traffic via internal queues in case the SAP queue in subscriber environment is disabled (**sub-sla-mgmt**>**single-sub-parameters**>**profiled-traffic-only**); this is applicable only to 1:1 subscriber model
- a possible necessity for HQoS Adjustment as multicast traffic is flowing outside of the subscriber queues
- de-coupling of the multicast forwarding statistics from the overall subscriber forwarding statistics obtained via subscriber specific **show** commands

8.4.1.1.2 IPoE 1:1 Model (Subscriber per VLAN/SAP) — No IGMP/MLD in AN

The AN is not IGMP/MLD aware, all replications are performed in the BNG. From the BNG perspective this deployment model has the following characteristics:

- IGMP/MLD states are kept per hosts and SAPs. Each host can be registered to more than one group.

- MLD utilize link-local as source address, which makes it difficult to associate with the originating host. MLD (S,G) are always first associated with a IPv6 WAN host if available. If this WAN host terminate its IPv6 session (via lease expiry, session terminate, and so on), the (S,G) is then associated to any remaining WAN host. Only when there are no more WAN hosts available is the (S,G) associated to the PD host, if any.
- IGMP/MLD Joins are accepted only from the active subscriber hosts as dictated by antispoofing.
- IGMP/MLD statistics can be displayed per host or per group.
- Multicast traffic for the subscriber is forwarded through the egress SAP queue. In case that the SAP queue is disabled (profiled-traffic-only command), multicast traffic flows via internal queues outside of the subscriber context.
- A single copy of any multicast stream is generated per SAP. This can be viewed as replication per unique multicast group per SAP, rather than the replication per host. In other words, the number of multicast streams on this SAP is equal to the number of unique groups across all hosts on this SAP (subscriber).
- Traffic statistics are kept per the SAP queue. Consequently multicast traffic stats are shown outside of the subscriber context.
- HQoS Adjustment might be necessary.
- Traffic cannot be explicitly classified (forwarding classes and queue mappings) inside of the subscriber queues.
- Redirection to the common multicast VLAN (or Layer 3 interface) is supported.
- Multicast streams have multicast destination MAC.
- MLD utilize link-local as source address, which makes it difficult to associate with the originating host. MLD (s,g) is associated with the IPv6 host. Therefore, if any WAN host or PD host end their IPv6 session (via lease expire, and so on), the (s,g) is associated with the remaining host address/prefix. The (s,g) is be delivered to the subscriber as long as a IPv6 address or prefix remains.

  This model is shown in Figure 55.

Figure 55 **1:1 Model**

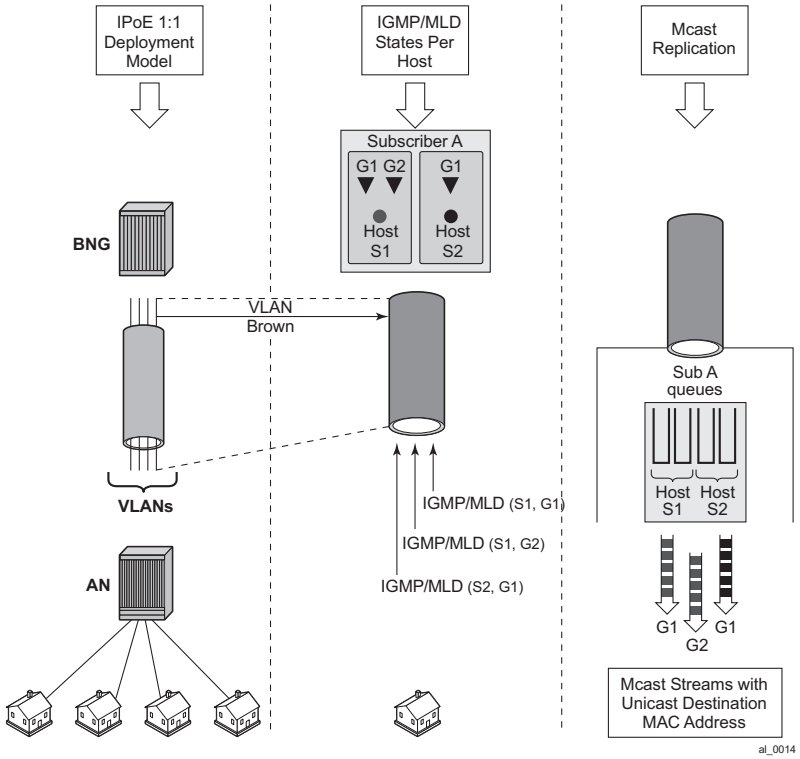

8.4.1.1.3 IPoE N:1 Model (Service per VLAN/SAP) — IGMP/MLD Snooping in the AN

The AN is IGMP/MLD aware and is participating in multicast replication. From the BNG perspective this deployment model has the following characteristics:

- IGMP/MLD states are kept per hosts and SAPs. Each host can be registered to more than one group.
- IGMP/MLD Joins are accepted only from the active subscriber hosts as dictated by antispoofing.
- IGMP/MLD statistics are displayed per host, per group or per subscriber.
- Multicast traffic for all subscribers on this SAP is forwarded through the egress SAP queues.

- A single copy of any multicast stream is generated per SAP. This can be viewed as the replication per unique multicast group per SAP, rather than the replication per host or subscriber. In other words, the number of multicast streams on this SAP is equal to the number of unique groups across all hosts and subscribers on this SAP.
- The AN receives a single multicast stream and based on its own (AN) IGMP/MLD snooping information, it replicates the mcast stream to the appropriate subscribers.
- Traffic statistics are kept per the SAP queue. Consequently multicast traffic stats are shown on a per SAP basis (aggregate of all subscribers on this SAP).
- Traffic cannot be explicitly classified (forwarding classes and queue mappings) inside of the subscriber queues.
- Redirection to the common multicast VLAN is supported.
- Multicast streams have multicast destination MAC.
- IGMP Joins are accepted (src IP address) only for the sub hosts that are already created in the system. IGMP Joins coming from the hosts that are nonexistent in the system are rejected, unless this functionality is explicitly enabled by the sub-hosts-only command under the IGMP group-int CLI hierarchy level.
- MLD utilize link-local as source address, which makes it difficult to associate with the originating host. MLD (S,G) are always first associated with a IPv6 WAN host, if available.If this WAN host terminate its IPv6 session (via lease expiry, session terminate, etc), the (S,G) is then associated to any remaining WAN host. Only when there are no more WAN hosts available, is the (S,G) associated to the PD host, if any.
- MLD join are only accepted if it matches the subscriber host link local address. MLD Joins coming from the hosts that are nonexistent in the system is rejected, unless this functionality is explicitly enabled by the sub-hosts-only command under the MLD group-int CLI hierarchy level.

  This model is shown in Figure 56.

Figure 56 **N:1 Model - AN in IGMP Snooping Mode**

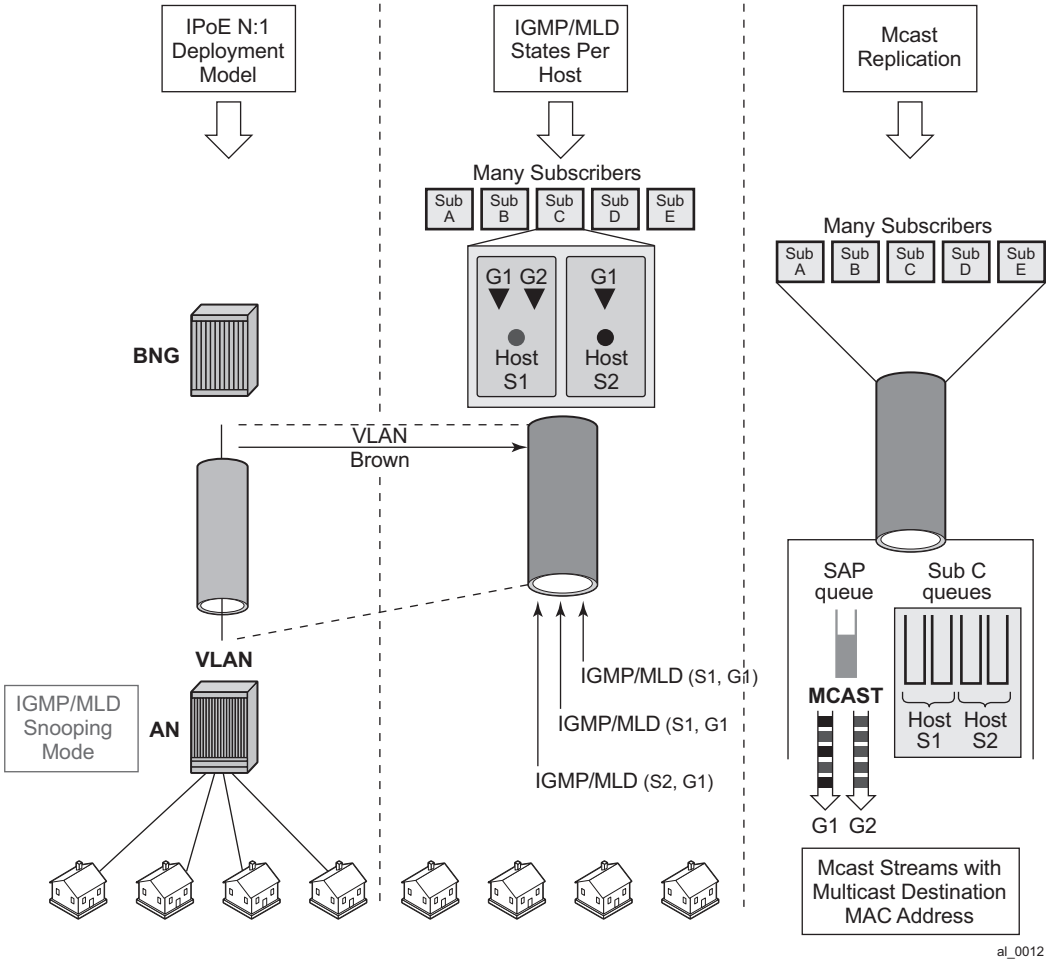

8.4.1.1.4 IPoE N:1 Model (Service per VLAN/SAP) — IGMP/MLD Proxy in the AN

The AN is configured as IGMP/MLD Proxy node and is participating in downstream multicast replication.

For IPv4, IGMP messages from multiple sources (subscribers hosts) for the same multicast group are consolidated in the AN into a single IGMP messages. This single IGMP message has the source IP address of the AN.

For IPv6, MLD messages from multiple sources (subscribers hosts) for the same multicast group are consolidated in the AN into a single MLD messages. This single MLD message has the link-local IP address of the AN.

From the BNG perspective this deployment model has the following characteristics:

- Subscriber IGMP/MLD states are maintained in the AN.
- IGMP Joins are accepted from the source IP address that is different from any of the subscriber's IP addresses already existing in the BNG. This is controlled via an IGMP filter on a per group-interface level assuming that the IGMP processing for subscriber hosts is disabled with the no **sub-hosts-only** command under the router/service **vprn>igmp>group-interface** CLI hierarchy. In this case all IGMP messages that cannot be related to existing hosts are treated in the context of the SAP while IGMP messages from the existing hosts are treated in the context of the subscriber hosts.
- MLD Joins are only accepted if the link-local address matches the subscriber' link local address. To allow processing of foreign link-local address such as the AN link local address, the MLD processing for subscriber hosts should be disabled with the no sub-hosts-only command under the **router/service vprn>mld>group-interface** CLI hierarchy. In this case all MLD messages that cannot be related to existing hosts are treated in the context of the SAP while MLD messages from the existing hosts are treated in the context of the subscriber hosts.
- IGMP/MLD statistics can be displayed per group-interface.
- Multicast traffic for all subscribers on this SAP is forwarded through the egress SAP queue.
- A single copy of any multicast stream is generated per SAP.
- The AN receives a single multicast stream. Based on the IGMP/MLD proxy information, the AN replicates the mcast stream to the appropriate subscribers.
- Traffic statistics are maintained per SAP queue.
- HQoS Adjustment is not useful because the per host/subscriber IGMP/MLD granularity is lost. IGMP/MLD states are aggregated per AN.
- Traffic can be explicitly classified into a specific SAP queues via a QoS policy applied under the SAP.
- Multicast streams have multicast destination MAC.

In the following example, IGMPs from the source IP address <ip> is accepted even though there is no subscriber-host with that IP addresses present in the system. An IGMP state is created under the SAP context (service per vlan, or N:1 model) for the **group** *pref-definition*. All other IGMP messages originated from non-subscriber hosts is rejected. IGMP messages for subscriber hosts are processed according to the IGMP policy applied to each subscriber host.

```
configure
     service vprn <id>
          igmp
               group-interface <name>
                    import <policy-name>

configure
     router
          policy-options
               begin
                    prefix-list <pref-name>
                         prefix <pref-definition>

                    policy-statement proxy-policy
                 entry 1
                    from
                          group-address <pref-name>
                            source-address <ip>
                            protocol igmp
                        exit
                   action accept
                   exit
                      exit
                      default-action reject
```

This functionality (accepting IGMP from non-subscriber hosts) can be disabled with the following flag.

```
configure
    service vprn <id>
        igmp
            group-interface <name>
    sub-host-only
```

In this case, only per host IGMP processing is allowed.

In the following example, MLDs with foreign link-local-address is accepted even though there is no subscriber-host with that link local addresses present in the system. An MLD state is created under the SAP context (service per vlan, or N:1 model) for the **group** *pref-definition*. All other MLD messages originated from non-subscriber hosts will be rejected. MLD messages for subscriber hosts are processed according to the IGMP policy applied to each subscriber host.

```
configure
     service vprn <id>
          mld
               group-interface <name>
                    import <policy-name>

configure
     router
          policy-options
               begin
```

```
prefix-list <pref-name>
     prefix <pref-definition>

policy-statement proxy-policy
entry 1
    from
          group-address <pref-name>
        source-address <ip>
        protocol igmp
    exit
   action accept
   exit
  exit
  default-action reject
```

This functionality (accepting MLD from non-subscriber hosts) can be disabled with the following flag.

```
configure
    service vprn <id>
        mld
            group-interface <name>
                sub-host-only
```

In this case, only per host MLD processing is allowed.

This model is shown in Figure 57.

Figure 57 **N:1 Model - AN in Proxy Mode**

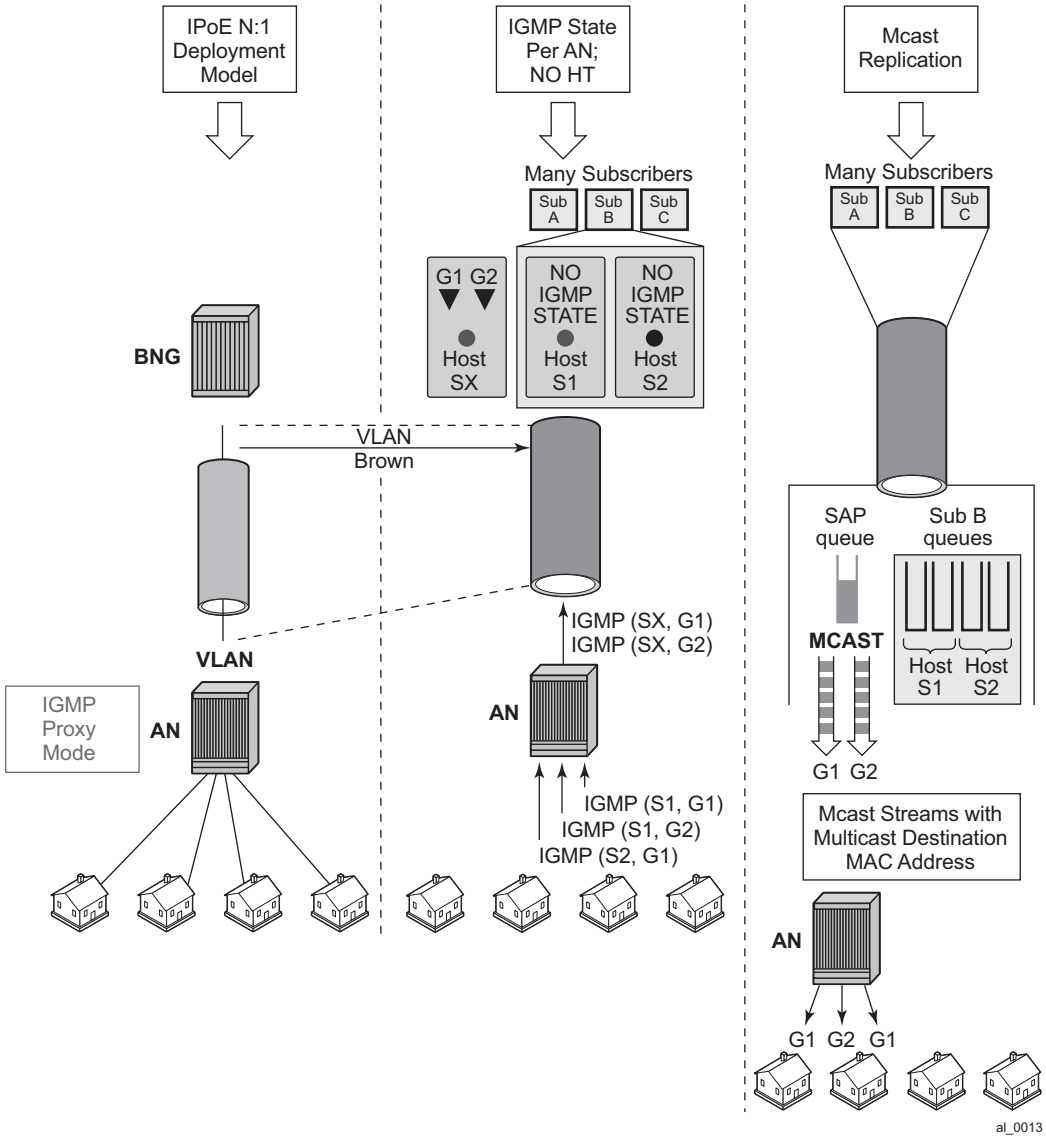

8.4.1.2 Per Subscriber Host Replication Mode

In this mode a multicast stream is transmitted per subscriber hosts for each registered multicast group (channel). As a result, multiple copies of the same multicast stream destined to different destinations can be transmitted over the same SAP. In this case traffic flows within the subscriber queues and consequently it is accounted in HQoS. As a result, HQoS Adjustment is not needed. Each copy of the same multicast stream have a unique unicast destination MAC addresses. The per host unicast MAC destination addresses are necessary to differentiate multiple copies between different receivers on the same SAP.

Per host replication mode can be enabled on a subscriber basis with the **per-host-replication** command in the **config>subscr-mgmt>igmp-policy** context.

For IPv6, the command is in the **config>subscr-mgmt>mld-policy** context.

8.4.1.2.1 IPoE 1:1 Model (Subscriber per VLAN/SAP) — No IGMP/MLD in AN

The AN is not IGMP/MLD aware and multicast replication is performed in the BNG. Multicast streams are sent directly to the hosts using their unicast MAC addresses. HQoS adjustment is not needed as multicast traffic is flowing through subscriber queues. From the BNG perspective this deployment model has the following characteristics:

- IGMP/MLD states are kept per hosts. Each host can be registered to multiple IGMP/MLD groups.
- MLD utilize link-local as source address, which makes it difficult to associate with the originating host. MLD (S,G) are always first associated with a IPv6 WAN host if available. If this WAN host terminate its IPv6 session (via lease expiry, session terminate, etc), the (S,G) is then associated to any remaining WAN host. Only when there are no more WAN hosts available, will the (S,G) be associated to the PD host, if any.
- IGMP/MLD Joins are accepted only from the active subscriber hosts. In other words antispoofing is in effect for IGMP/MLD messages.
- IGMP/MLD statistics can be displayed per host, per group or per subscriber.
- Multicast traffic is forwarded through subscriber queues using unicast destination MAC address of the destination host.
- Multiple copies of the same multicast stream can be generated per SAP. The number of copies depends on the number of hosts on the SAP that are registered to the same multicast group (channel). In other words, the number of multicast streams on the SAP is equal to the number of groups registered across all hosts on this SAP.

- Traffic statistics are kept per the host queue. In case that multicast statistics need to be separated from unicast, the multicast traffic should be classified in a subscriber separate queue.
- HQoS Adjustment is not needed as traffic is flowing within the subscriber queues and is automatically accounted in HQoS.
- Multicast traffic can be explicitly classified into forwarding classes and consequently directed into desired queues.
- MCAC is supported.
- profiled-traffic-only mode defined under sub-sla-mgmt is supported. This mode (profiled-traffic-only) is used to save the number of queues in 1:1 model (sub-sla-mgmt-> no multi-sub-SAP) by preventing the creation of the SAP queues. Since multicast traffic is not using the SAP queue, enabling this feature will not have any effect on the multicast operation.

  This model is shown in Figure 58.

Figure 58 **1:1 Model**

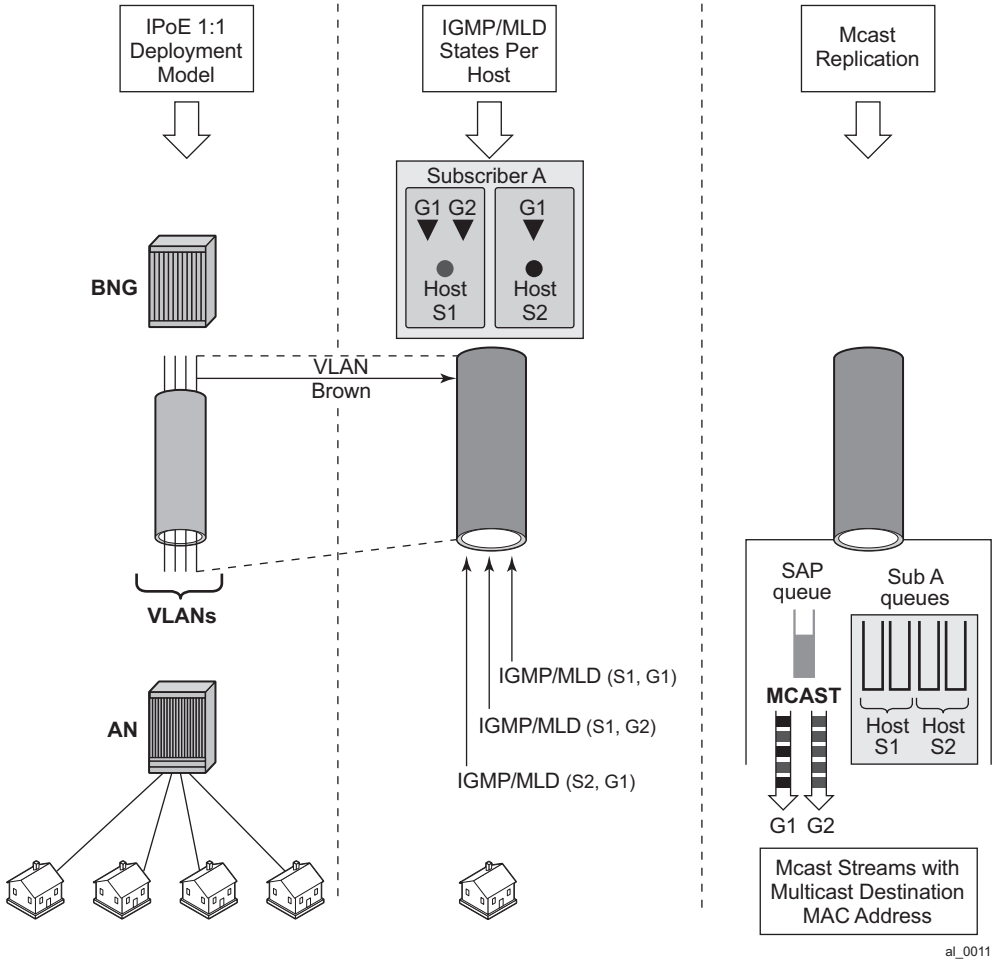

8.4.1.2.2 IPoE N:1 Model (Service per VLAN/SAP) — No IGMP/MLD in the AN

The AN is not IGMP/MLD aware and is not participating in multicast replication. From the BNG perspective this deployment model has the following characteristics:

- IGMP/MLD states are kept per hosts. Each host can be registered to multiple multicast groups.

- MLD utilize link-local as source address, which makes it difficult to associate with the originating host. MLD (S,G) are always first associated with a IPv6 WAN host if available. If this WAN host terminate its IPv6 session (via lease expiry, session terminate, etc), the (S,G) is then associated to any remaining WAN host. Only when there are no more WAN hosts available, will the (S,G) be associated to the PD host, if any.
- IGMP/MLD Joins will be accepted only from the active subscriber hosts, subject to antispoofing.
- IGMP/MLD statistics can be displayed per host, per group or per subscriber.
- Multicast traffic is forwarded through subscriber queues using unicast destination MAC address of the destination host.
- Multiple copies of the same multicast stream can be generated per SAP. The number of copies depends on the number of hosts on the SAP that are registered to the same multicast group (channel). In other words, the number of multicast streams on the SAP is equal to the number of groups registered across all hosts on this SAP.
- Traffic statistics are kept per the host queue. In case that multicast statistics need to be separated from unicast, the multicast traffic should be classified in a separate subscriber queue.
- HQoS Adjustment is not needed as traffic is flowing within the subscriber queues and is automatically accounted in HQoS.
- Multicast traffic can be explicitly classified into forwarding classes and consequently directed into desired queues.
- MCAC is supported.

  This model is shown in Figure 59.

Figure 59 **N:1 Model — No IGMP/MLD in the AN**

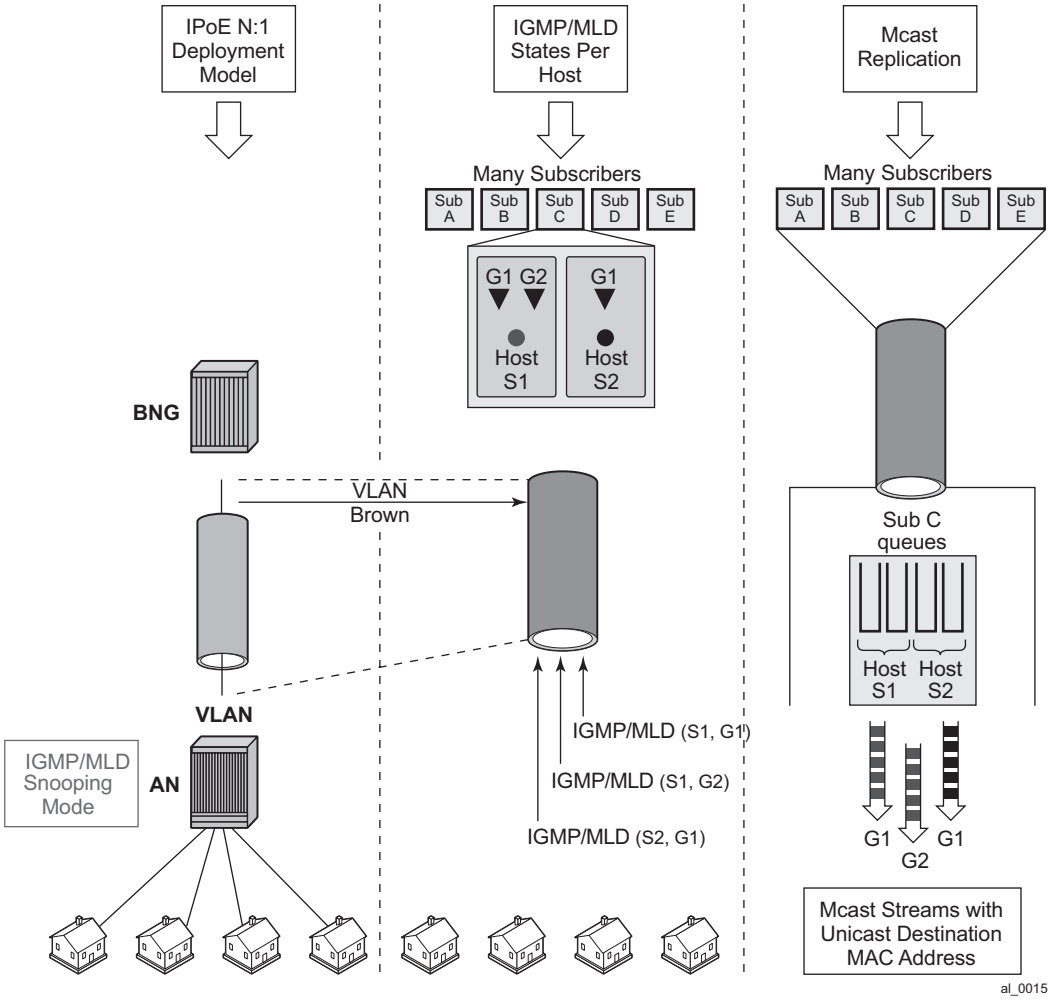

8.4.2 Multicast Over PPPoE

In a PPPoE environment, multicast replication is performed per session (host) regardless of whether those sessions are shared per SAP or they reside on individual SAPs. This is due to the point-to-point nature of PPPoE sessions. HQoS adjustment is not needed as multicast is part of the PPPoE session traffic that is flowing via subscriber queues. Multicast packets are sent with unicast MAC address to each CPE. PPP protocol field is set to IP and the destination IP address is the multicast group address for each unique session ID.

This model is shown in Figure 60.

Figure 60 **Multicast IPv4 Address and Unicast MAC Address in PPPoE Subscriber Multicast**

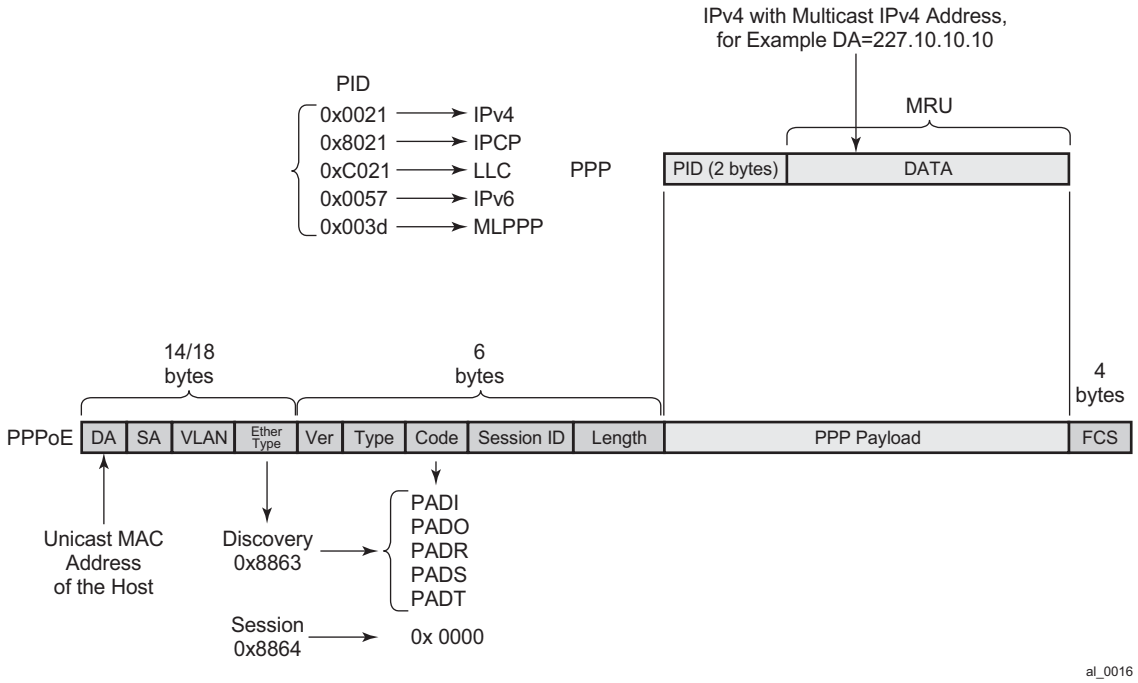

Note that in a SRRP setup, the SRRP master utilizes the SRRP MAC address as the source MAC for both multicast and multicast control protocol packets. This is only applicable to PPPoE hosts.

8.4.3 IGMP Flooding Containment

The query function in IGMP can cause some unintended flooding in N:1 IPoE deployment model with AN in the IGMP snooping mode. By maintaining IGMP session states per host, it is assumed that the IGMP interaction between multicast receivers and the BNG will be on a one-to-one basis. Upon arrival of an IGMP leave from a host for a specific multicast group, the IGMP querier would normally multicast a group-specific query (fast-leave). In N:1 model with SAP replication mode enabled, the 7750 SR and 7450 ESS sends a group-specific query (fast-leave) only when it receives the IGMP leave message for the last group shared amongst all subscribers on this SAP.

8.4.4 IGMP/MLD Timers

IGMP/MLD timers are maintained under the following hierarchy:

IPv4:

config>router>igmp

config>service vprn>igmp

IPv6:

config>router>mld

config>service vprn>mld

The IGMP/MLD timers are controlled on a per routing instance (VRF or GRT) level.

The timer values are used to:

- Determine the interval at which queries are transmitted (query-interval).
- To determine the amount of time after which a join will time out.

However, the timers can be different for hosts and redirected interface in case that redirection between VRFs is enabled.

8.4.5 IGMP/MLD Query Intervals

IGMP/MLD query related intervals (query-interval, query-last-member-interval, query-response-interval, robust-count) are configured on a global router/vprn IGMP/MLD level. They are used to determine the IGMP/MLD timeout states and the rates at which queries are transmitted.

In case of redirection, the subscriber-host IGMP/MLD state will determine the IGMP/MLD state on the redirected interface, assuming that IGMP/MLD messages are not directly received on the redirected interface (for example from the AN performing IGMP/MLD forking). For example if the redirected interface is not receiving IGMP/MLD messages from the downstream node, then the IGMP/MLD state under redirected interface will be removed simultaneously with the removal of the IGMP/MLD state for the subscriber host (due to leave or a timeout).

In case that the redirected interface is receiving IGMP/MLD message directly from the downstream node, the IGMP/MLD states on that redirected interface will be driven by those direct IGMP/MLD messages.

For example, an IGMP/MLD host in VRF1 has an expiry time of 60 seconds and the expiry time defined under the VRF2 where multicast traffic is redirected is set to 90 seconds. The IGMP/MLD state will time out for the host in VRF1 after 60s, and if no host has joined the same multicast group in VRF2 (where redirected interface resides), the IGMP state will be removed there too.

If a join was received directly on the redirection interface in VRF2, the IGMP/MLD state for that group will be maintained for 90s, regardless of the IGMP/MLD state for the same group in VRF1.

8.4.6 HQoS Adjustment

HQoS Adjustment is required in the scenarios where subscriber multicast traffic flow is disassociated from subscriber queues. In other words, the unicast traffic for the subscriber is flowing through the subscriber queues while at the same time multicast traffic for the same subscriber is explicitly (through redirection) or implicitly (per-sap replication mode) redirected through a separate non-subscriber queue. In this case HQoS Adjustment can be deployed where preconfigured multicast bandwidth per channel is artificially included in HQoS. For example, bandwidth consumption per multicast group must be known in advance and configured within the 7450 ESS and 7750 SR. By keeping the IGMP state per host, the bandwidth for the multicast group (channel) to which the host is registered is known and is deducted as consumed from the aggregate subscriber bandwidth.

The multicast bandwidth per channel must be known (this is always an approximation) and provisioned in the BNG node in advance.

In PPPoE and in IPoE per host replication environment, HQoS Adjustment is not needed as multicast traffic is unicasted to each subscriber and therefore is flowing through subscriber queues.

For HQoS Adjustment, the channel bandwidth definition and association with an interface is the same as in the MCAC case. This is a departure from the legacy HT channel bandwidth definition which is done via a multicast-info-policy.

Example of HQoS adjustment:

Channel definition:

```
configure
    router
        mcac
            policy <name>
                <channel definition>
```

Channel bandwidth definition policy can be applied under:

- group-interface

```
configure
    service vprn <id>
        igmp
            group-interface <grp-if-name>
                mcac
                    policy <mcac-policy-name>
```

- plain interface

```
configure
    router/service vprn
        igmp
            interface <name>
                mcac
                    policy <mcac-policy-name>
```

- retailer group-interface:

```
configure
    service vprn <id>
        igmp
            group-interface fwd-service <svc-id> <grp-if-name>
                mcac
  policy <mcac-policy-name>
```

Enabling HQoS adjustment:

```
configure
    subscriber-management
        igmp-policy <name>
            egress-rate-modify [egress-aggregate-rate-limit | scheduler <name>]
```

Applying HQoS adjustment to the subscriber:

```
configure
    subscriber-management
        sub-profile <name>
            igmp-policy <name>
```

In order to activate HQoS adjustment on the subscriber level, the sub-mcac-policy must be enabled under the subscriber via the following CLI:

```
configure
    subscriber-management
        sub-mcac-policy <pol-name>
            no shutdown

configure
    subscriber-management
        sub-profile <name>
            sub-mcac-policy <pol-name>
```

The adjusted bandwidth during operation can be verified with the following commands (depending whether agg-rate-limit or scheduler-policy is used):

```
*B:BNG-1# show service active-subscribers subscriber "sub-1" detail
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber sub-1
-------------------------------------------------------------------------------
I. Sched. Policy : up-silver
E. Sched. Policy : N/A                           E. Agg Rate Limit: 4000
I. Policer Ctrl. : N/A
E. Policer Ctrl. : N/A
Q Frame-Based Ac*: Disabled
Acct. Policy     : N/A                           Collect Stats    : Enabled
Rad. Acct. Pol.  : sub-1-acct
Dupl. Acct. Pol. : N/A
ANCP Pol.        : N/A
HostTrk Pol.     : N/A
IGMP Policy      : sub-1-IGMP-Pol
Sub. MCAC Policy : sub-1-MCAC
NAT Policy       : N/A
Def. Encap Offset: none                          Encap Offset Mode: none
Avg Frame Size   : N/A
Preference       : 5
Sub. ANCP-String : "sub-1"
Sub. Int Dest Id : ""
Igmp Rate Adj    : -2000
RADIUS Rate-Limit: N/A
Oper-Rate-Limit  : 2000
...
-------------------------------------------------------------------------------
*B:BNG-1#
```

Consider a different example with a scheduler instead of agg-rate-limit:

```
*A:Dut-C>config>subscr-mgmt>sub-prof# info
----------------------------------------------
            igmp-policy "pol1"
            sub-mcac-policy "smp"
            egress
                scheduler-policy "h1"
                    scheduler "t2" rate 30000
                exit
            exit
----------------------------------------------

*A:Dut-C>config>subscr-mgmt>igmp-policy# info
----------------------------------------------
            egress-rate-modify scheduler "t2"
            redirection-policy "mc_redir1"
----------------------------------------------
```

Assume that the subscriber joins now a new channel with bandwidth of 1 Mb/s (1000 kb/s).

```
A:Dut-C>config>subscr-mgmt>sub-prof>egr>sched# show qos scheduler-
hierarchy subscriber "sub_1" detail
===============================================================================
Scheduler Hierarchy - Subscriber sub_1
===============================================================================
Ingress Scheduler Policy:
Egress Scheduler Policy : h1
-------------------------------------------------------------------------------
Legend :
(*) real-time dynamic value
(w) Wire rates
B   Bytes
-------------------------------------------------------------------------------
Root (Ing)
|
No Active Members Found on slot 1

Root (Egr)
| slot(1)
|--(S) : t1
|   |    AdminPIR:90000      AdminCIR:10000
|   |
|   |
|   |    [Within CIR Level 0 Weight 0]
|   |    Assigned:0          Offered:0
|   |    Consumed:0
|   |
|   |    [Above CIR Level 0 Weight 0]
|   |    Assigned:0          Offered:0
|   |    Consumed:0
|   |    TotalConsumed:0
|   |    OperPIR:90000
|   |
|   |    [As Parent]
|   |    Rate:90000
|   |    ConsumedByChildren:0
|   |
|   |
|   |--(S) : t2
|   |   |    AdminPIR:29000      AdminCIR:10000(sum)       <==== bw 1000 from igmp s
ubstracted
|   |   |
|   |   |
|   |   |    [Within CIR Level 0 Weight 1]
|   |   |    Assigned:10000      Offered:0
|   |   |    Consumed:0
|   |   |
|   |   |    [Above CIR Level 1 Weight 1]
|   |   |    Assigned:29000      Offered:0                 <==== bw 1000 from igmp s
ubstracted
|   |   |    Consumed:0
|   |   |
|   |   |
|   |   |    TotalConsumed:0
|   |   |    OperPIR:29000                                 <==== bw 1000 from igmp s
ubstracted
|   |   |
|   |   |    [As Parent]
|   |   |    Rate:29000                                     <==== bw 1000 from igmp
```

```
substracted
|   |   |    ConsumedByChildren:0
|   |   |
|   |   |
|   |   |--(S) : t3
|   |   |   |    AdminPIR:70000      AdminCIR:10000
|   |   |   |
|   |   |   |
|   |   |   |    [Within CIR Level 0 Weight 1]
|   |   |   |    Assigned:10000      Offered:0
|   |   |   |    Consumed:0
|   |   |   |
|   |   |   |    [Above CIR Level 1 Weight 1]
|   |   |   |    Assigned:29000      Offered:0
|   |   |   |    Consumed:0
|   |   |   |
|   |   |   |
|   |   |   |    TotalConsumed:0
|   |   |   |    OperPIR:29000
|   |   |   |
|   |   |   |    [As Parent]
|   |   |   |    Rate:29000
|   |   |   |    ConsumedByChildren:0
|   |   |   |

*A:Dut-C>config>subscr-mgmt>igmp-policy# show service active-subscribers sub-mcac
===============================================================================
Active Subscribers Sub-MCAC
===============================================================================
Subscriber                          : sub_1
MCAC-policy                         : smp (inService)
In use mandatory bandwidth          : 1000
In use optional bandwidth           : 0
Available mandatory bandwidth       : 1147482647
Available optional bandwidth        : 1000000000
-------------------------------------------------------------------------------
Subscriber                          : sub_2
MCAC-policy                         : smp (inService)
In use mandatory bandwidth          : 0
In use optional bandwidth           : 0
Available mandatory bandwidth       : 1147483647
Available optional bandwidth        : 1000000000
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Number of Subscribers : 2
===============================================================================
*A:Dut-C#

*A:Dut-C# show service active-subscribers subscriber "sub_1" detail
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber sub_1 (1)
-------------------------------------------------------------------------------
I. Sched. Policy : N/A
E. Sched. Policy : h1                               E. Agg Rate Limit: Max
I. Policer Ctrl. : N/A
E. Policer Ctrl. : N/A
```

```
Q Frame-Based Ac*: Disabled
Acct. Policy     : N/A                              Collect Stats    : Disabled
Rad. Acct. Pol.  : N/A
Dupl. Acct. Pol. : N/A
ANCP Pol.        : N/A
HostTrk Pol.     : N/A
IGMP Policy      : pol1
Sub. MCAC Policy : smp
NAT Policy       : N/A
Def. Encap Offset: none                             Encap Offset Mode: none
Avg Frame Size   : N/A
Preference       : 5
Sub. ANCP-String : "sub_1"
Sub. Int Dest Id : ""
Igmp Rate Adj    : N/A
RADIUS Rate-Limit: N/A
Oper-Rate-Limit  : Maximum
...
===============================================================================
*A:Dut-C#


*A:Dut-C# show subscriber-mgmt igmp-policy "pol1"
===============================================================================
IGMP Policy pol1
===============================================================================
Import Policy                       :
Admin Version                       : 3
Num Subscribers                     : 2
Host Max Group                      : No Limit
Host Max Sources                    : No Limit
Fast Leave                          : yes
Redirection Policy                  : mc_redir1
Per Host Replication                : no
Egress Rate Modify                  : "t2"
Mcast Reporting Destination Name    :
Mcast Reporting Admin State         : Disabled
===============================================================================
*A:Dut-C#
```

8.4.6.1 Host Tracking (HT) Considerations

HT is a light version of HQoS Adjustment feature. The use of HQoS Adjustment functionality in place of HT is strongly encouraged.

When HT is enabled, the AN forks off (duplicates) the IGMP messages on the common mcast SAP to the subscriber SAP. IGMP states are not fully maintained per sub-host in the BNG, instead they are only tracked (less overhead) for bandwidth adjustment purposes.

Example of HT

```
Channel Definition:
configure
    mcast-management
        multicast-info-policy <name>
            <channel to b/w mapping definition>
```

Applying channel definition policy on a router/VPRN global level:

```
configure>router>multicast-info-policy <name>
configure>service>vrpn>multicast-info-policy <name>
```

Defining the rate object on which HT is applied:

```
configure
    subscriber-management
        host-tracking-policy <name>
            egress-rate-modify [agg-rate-limit | scheduler <sch-name>]
```

Applying the HT to the subscriber:

```
configure
    subscriber-management
        sub-profile <name>
            host-tracking-policy <name>  => mutually exclusive with igmp-policy
```

8.4.6.2 HQoS Adjust Per Vport

HQoS adjust per Vport can be used in environments where Vport represents a physical medium over which traffic for multiple subscribers is shared. Typical example of this scenario is shown in Figure 61. Multicast traffic within the router is taking a separate path from unicast traffic, only for the two traffic flows to merge later in the PON (represented by a Vport) and ONT (represented by the subscriber in the 7450 ESS and 7750 SR).

Figure 61 **HQoS Adjustment per Subscriber and Vport**

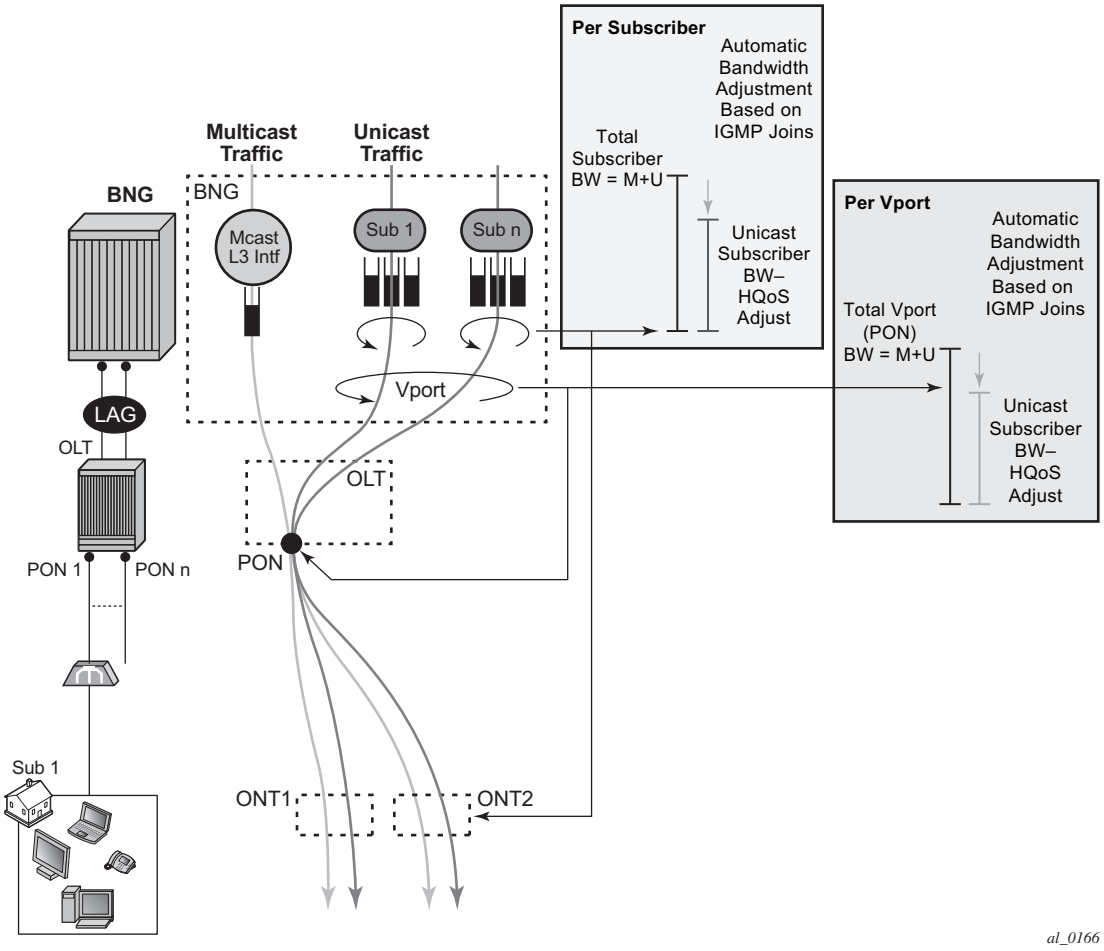

A single copy of each channel is replicated on the PON as long as there is at least one subscriber on that PON interested in this channel (has joined the IGMP/MLD group).

The 7450 ESS and 7750 SR monitors IGMP/MLD Joins at the subscriber level and consequently the channel bandwidth is subtracted from the current Vport rate limit only in the case that this is the first channel flowing through the corresponding PON. Otherwise, the Vport bandwidth is not modified. Similarly, when the channel is removed from the last subscriber on the PON, the channel bandwidth is returned to the VPort.

Association between the Vport and the subscriber is performed via inter-destination-string or svlan during the subscriber setup phase. Inter-destination-string can be obtained either via RADIUS or LUDB. In case that the association between the Vport and the subscriber is performed based on the svlan (as specified in sub-sla-mgmt under the sap/msap), then the destination string under the Vport must be a number matching the svlan.

The mcac-policy (channel definition bandwidth) can be applied on the group interface under which the subscribers are instantiated or in case of redirection under the redirected-interface.

In a LAG environment, the Vport instance is instantiated per member LAG link on the IOM. For accurate bandwidth control, it is prerequisite for this feature that subscriber traffic hashing is performed per Vport.

The CLI structure is as follows.

```
configure
    port <port-id>
        ethernet
            access
                egress
                    vport <name>
                        egress-rate-modify
                        agg-rate
                        host-match <destination-string>
                        port-scheduler-policy <port-scheduler-policy-name>

configure
    port <port-id>
        sonnet-sdh
            path [<sonnet-sdh-index>]
                access
                    egress
                        vport <name>
                            egress-rate-modify
                            agg-rate
                            host-match <destination-string>
                            port-scheduler-policy <port-scheduler-policy-name>
```

The Vport rate that is affected by this functionality depends on the configuration:

- When the agg-rate-limit within the Vport is configured, its value is modified based on the IGMP activity associated with the subscriber under this Vport.
- When the port-scheduler-policy within the Vport is referenced, the *max-rate* defined in the corresponding port scheduler policy is modified based on the IGMP activity associated with the subscriber under this Vport.

Note: HQoS adjust is not supported when a scheduler policy is configured under the VPORT.

The Vport rates can be displayed with the following two commands:

show port 1/1/5 vport *name*

qos scheduler-hierarchy port *port-id* **vport** *vport-name*

As an example:

```
*A:system-1# show port 1/1/7 vport
===============================================================================
Port 1/1/7 Access Egress vport
===============================================================================
VPort Name    : isam1
Description   : (Not Specified)
Sched Policy  : 1
Rate Limit    : Max
Rate Modify   : enabled
Modify delta  : -14000
```

In this case, the configured Vport aggregate-rate-limit max value has been reduced by 14 Mb/s.

Similarly, if the Vport had a port-scheduling-policy applied, the max-rate value configured in the port-scheduling-policy would have been modified by the amount shown in the Modify delta output in the above command.

8.4.6.2.1 Multi-Chassis Redundancy

Modified Vport rate synchronization in multi-chassis environment relies on the synchronization of the subscriber IGMP/MLD states between the redundant nodes. Upon the switchover, the Vport rate on the newly active node is adjusted according to the current IGMP/MLD state of the subscribers associated with the Vport.

8.4.6.2.2 Scalability Considerations

It is assumed that the rate of the IGMP/MLD state change on the Vport level is substantially lower than on the subscriber level.

The reason for this is that the IGMP/MLD Join/Leaves are shared amongst subscribers on the same Vport (PON for example) and thus the IGMP/MLD state on the VPort level is changed only for the first IGMP/MLD Join per channel and the last IGMP/MLD leave per channel.

8.4.7 Redirection

Two levels of MCAC can be enabled simultaneously and in such case this is referred as Hierarchical MCAC (H-MCAC). In case that redirection is enabled, H-MCAC per subscriber and the redirected interface is supported. However, mcac per group-interface in this case is not supported. Channel definition policy for the subscriber and the redirected interface is in this case referenced under the **igmp->interface** (redirected interface) CLI or for IPv6 **mld->interface**.

In case that redirection is disabled, H-MCAC for both, the subscriber and the group-interface is supported. The channel definition policy is in this case configured under the **config>router>igmp>group interface** context or for IPv6 **config>router>mld>group interface**.

There are two options in multicast redirection. The first option is to redirect all subscriber multicast traffic to a dedicated redirect interface.

Example:

Defining redirection action:

```
configure
    router
        policy-options
            begin
            policy-statement <name>
                default-action accept
                multicast-redirection [fwd-service <svc id>] <interface name>
                exit
            exit
        exit
    exit
```

The second option is to redirect only specific multicast groups to the redirect interface while the remaining groups remains on the subscriber SAP. This is applicable for both IPv4 and IPv6. For IPv6 host-ip for a policy statement is not supported.

Example:

Defining redirection action:

```
configure
    router
        policy-options
            begin
            prefix-list <name>
                prefix <IPv4 multicast groups>
                prefix <IPv4 multicast groups
            exit
            policy-statement <name>
                entry 1
                    from
                        group-address <prefix-list name>
                    action accept
                    multicast-redirection [fwd-service <svc id>] <interface name>
                exit
            exit
    exit
```

Applying redirection to the subscriber for IGMP and MLD respectively.

```
    configure
        subscr-mgmt
            igmp-policy <name>
                redirection-policy <name>
                exit
            exit
            mld-policy <name>
                redirection-policy <name>
```

Redirection that cross-connects GRT and VPRN is not supported. Redirection can be only performed between interfaces in the GRT, or between the interfaces in any of the VPRN (cross connecting VPRNs is allowed).

Redirection is also supported in a wholesaler/retailer VPRN model where redirected Layer 3 interface resides in the retailer VPRN.

8.4.8 Hierarchical Multicast CAC (H-MCAC)

MCAC is supported on three levels:

- per subscriber
- per group-interface
- per redirected interface

Two levels of MCAC can be enabled simultaneously and in such case this is referred as Hierarchical MCAC (H-MCAC). In case that redirection is enabled, H-MCAC per subscriber and the redirected interface is supported. However, MCAC per group-interface in this case is not supported. Channel definition policy for the subscriber and the redirected interface is in this case referenced under the **config>router>igmp->interface** (redirected interface) CLI hierarchy or IPv6 **config>router>mld->interface**.

In case that redirection is disabled, H-MCAC for both, the subscriber and the group-interface is supported. The channel definition policy is in this case configured under the **config>router>igmp>group-interface** CLI hierarchy or IPv6 **config>router>mld>group- interface**.

Examples

Note: The same channel definition and association with interfaces is used for MCAC/H-MCAC and HQoS Adjustment.

Channel definition:

```
configure
    router
        mcac
            policy <mcac-pol-name>
                bundle <bundle-name>
                    bandwidth <kbps>
                    channel <start-address> <end-
address> bw <bw> [class {high|low}] [type {mandatory|optional}]
                    :
                    :
```

Channel bandwidth definition policy can be referenced under the:

- group-interface — This is used for subscribers when redirection is disabled.

```
configure
                service vprn <id>
            igmp/mld
                 group-interface <grp-if-name>
                    mcac
                        policy <mcac-policy-name>
                        policy <mcac-policy-name>
```

- plain interface

```
configure
    router
        igmp/mld
            interface <name>
                mcac
                    policy <mcac-policy-name>
```

```
configure
    service vprn <id>
        igmp/mld
            interface <if-name>
                mcac
                    unconstrained-bw <bandwidth> mandatory-bw <mandatory-bw>
```

- retailer's VPRN reference the group-interface in the wholesaler's VPRN

```
configure
    service vprn <id>
        igmp/mld
            group-interface fwd-service <svc-id> <grp-if-name>
                mcac
                    policy <mcac-policy-name>
```

Enabling MCAC per subscriber:

```
configure
    subscr-mgmt
        sub-mcac-policy <name>
            unconstrained-bw <bandwidth> mandatory-bw <mandatory-bw>

configure
    subscr-mgmt
        sub-profile <name>
            sub-mcac-policy <name>
```

Enabling MCAC per-group-interface

```
configure
        service vprn <id>
            igmp/mld
                group-interface <grp-if-name>
                    mcac
                        unconstrained-bw <bandwidth> mandatory-bw <mandatory-bw>
```

Enabling MCAC per redirected interface

```
configure
    router
        igmp/mld
            interface <if-name>
                mcac
                    unconstrained-bw <bandwidth> mandatory-bw <mandatory-bw>

configure
    service vprn <id>
        igmp/mld
            interface <if-name>
                mcac
                    unconstrained-bw <bandwidth> mandatory-bw <mandatory-bw>
```

8.4.8.1 MCAC Bundle Bandwidth Limit Considerations

In addition to multicast bandwidth limit that can be imposed on subscribers, group-interfaces or regular interfaces, there is another multicast bandwidth limit that can be imposed on a group of channels (channel bundle).

The MCAC policy, aside from the channel bandwidth definitions, could optionally contain this bandwidth cap for the group of channels:

```
config>router>mcac# info
----------------------------------------------
        policy "test"
            bundle "test" create
                bandwidth 100000
                channel 239.0.0.10 239.0.0.10 bw 10000 type mandatory
                channel 239.0.0.11 239.0.0.15 bw 5000 type mandatory
                channel 239.0.0.20 239.0.0.30 bw 5000 type optional
            exit
        exit
```

This can be used to prevent a single set of channels from monopolizing MCAC bandwidth allocated to the entire interface. The bandwidth of each individual bundle is capped to some value below the interface MCAC bandwidth limit, allowing each bundle to have its own share of the interface MCAC bandwidth.

In most cases, the bandwidth limit per bundle is not necessary to configure. The aggregate limit per all channels as defined under the subscriber or interface covers majority of scenarios. In case that one wants to explore the bundle bandwidth limits and how they affect MCAC behavior, the following information helps understanding this topic.

To further understand how various MCAC bandwidth limits are applied, one need to understand the concept of the mandatory bandwidth that is pre-allocated in the following way:

- Bandwidth of each mandatory channel in a bundle is pre-allocated. The artifacts of this are:
  - The total mandatory bandwidth in the bundle cannot exceed the bundle cap. For the sake of deterministic behavior, the configured bandwidth of each mandatory channel in the bundle is counted towards the total mandatory bandwidth only once. This means that only one replication of each mandatory channel is assumed. This is normal behavior on a regular interface with a single SAP under it. More than one replication of the same channel per regular interface (or SAP) would lead to packet duplication.

  - Optional (non-mandatory) channels can use only the difference in bandwidth between the bundle cap and total pre-allocated mandatory bandwidth. They cannot use more bandwidth than that even if the total pre-allocated mandatory bandwidth is not used up (mandatory channels are not being replicated).

- Mandatory bandwidth under the interface is pre-allocated and subtracted from the unconstrained bandwidth. In the configuration example below, 2 Mb/s is pre-allocated (guaranteed for mandatory channels) and the remaining 8 Mb/s can be used by the optional channels on a first come first serve basis.

```
config>router>igmp# info
----------------------------------------------
            interface "ge-1/1/1"
                mcac
                    unconstrained-bw 10000 mandatory-bw 2000
                exit
            exit
```

The bundle bandwidth limit poses a problem when the MCAC policy is applied under the group interface. The reason is that the group interface represents the aggregation point for the subscribers and their bandwidth. As such it is natural that the any aggregated bandwidth limit under the group interface be larger than the bandwidth limit applied to any individual subscriber under it. Since the MCAC policy, along with the bundle bandwidth limit, is inherited by all subscribers under the group-interface, the exhaustion of the bundle bandwidth limit under the group interface coincides with the exhaustion of the bundle bandwidth limit of any individual subscriber. This results in a single subscriber starving out of multicast bandwidth the remaining subscribers under the same group interface. While it is perfectly acceptable for the subscribers to inherit the multicast channel definition from the group-interface, for the above reasons it is not acceptable that the subscriber inherit the bandwidth cap from the group-interface.

To remedy this situation, the MCAC bandwidth limits are independently configured under the group interface level (aggregated level) and the subscriber level via the **unconstrained-bw** *kbps* **mandatory-bw** *kbps* command. The undesired bundle bandwidth cap in the MCAC policy is ignored under the group-interface and under the subscriber. However, the bundle bandwidth cap is applied automatically to each SAP under the group interface. A SAP is a natural place for a bundle bandwidth limit since each channel on a SAP can be replicated only once and therefore the amount of pre-allocated mandatory bandwidth can be pre-calculated. This is obviously not the case for the group interface where single channel can be replicated multiple times (one per each SAP under the grp-if). Similarly, the same channel can be replicated multiple times for the same subscriber in per-host replication mode. Only subscribers in per-sap replication mode warrants a single replication per channel. Therefore, if bundle cap is configured, it is applied to limit the bandwidth of the bundle that is applied to a subscriber in a per-sap replication mode.

Figure 62 depicts MCAC related inheritances and MCAC bandwidth allocation model in per-sap replication mode. The MCAC policy is applied to the group interface and inherited by each subscriber as well as each SAP under the same group interface. However, the bundle bandwidth limit in the MCAC policy is ignored on the group-interface and under the subscriber (denoted by the X in the figure). The bundle limit is applied only to each sap under the group-interface.

Overall (non-bundle) MCAC bandwidth limits are independently applied to the group-interface and the subscribers. According to our example, 20 Mb/s of multicast bandwidth in total is allocated per group-interface. 6 Mb/s of the 20 Mb/s is allocated for mandatory channels. This leaves 14 Mb/s of multicast bandwidth for the optional channels combined served on a first come first serve basis. Each physical replication (multiple replications of the same channel can occur, one per each SAP), counts towards the respective group-interface bandwidth limits.

Similar logic applies to the subscriber MCAC bandwidth limits which are applied per sub-profile.

Finally, each SAP can optionally contain the bundle bandwidth limit.

Note: In a hierarchical MCAC fashion, if either of the bandwidth checks fails (SAP, sub or grp-if) the channel admission for the subscriber also fails.

In the example, six subscriber hosts watch the same channel but there are only three active replications (one per SAP). This would yield:

- 14 Mb/s of available multicast bandwidth under the group-interface. This bandwidth can be used for optional channels on a first come first serve basis. No reserved bandwidth is left.
- Subscriber A — 3 Mb/s is still reserved for mandatory channels and 5 MB/s is available for optional channels (first come first serve). All this assume that the SAP and the grp-if bandwidth checks pass.
- Subscriber B — 2 Mb/s is still reserved for mandatory channels and 6 Mb/s is available for optional channels (first come first serve). All this assume that the SAP and the grp-if bandwidth checks pass.
- Subscriber C — No reserved bandwidth for mandatory channels is left. 3 Mb/s is still left for optional channels. All this assume that the SAP and the grp-if bandwidth checks pass.
- SAPs — Considering that 2 Mb/s are currently replicating (ch A, each SAP can still accept 1 Mb/s of the mandatory bandwidth (channel B and 7 Mb/s of the remaining optional channels.

Figure 62 **MCAC Policy Inheritance in Per-SAP Replication Mode**

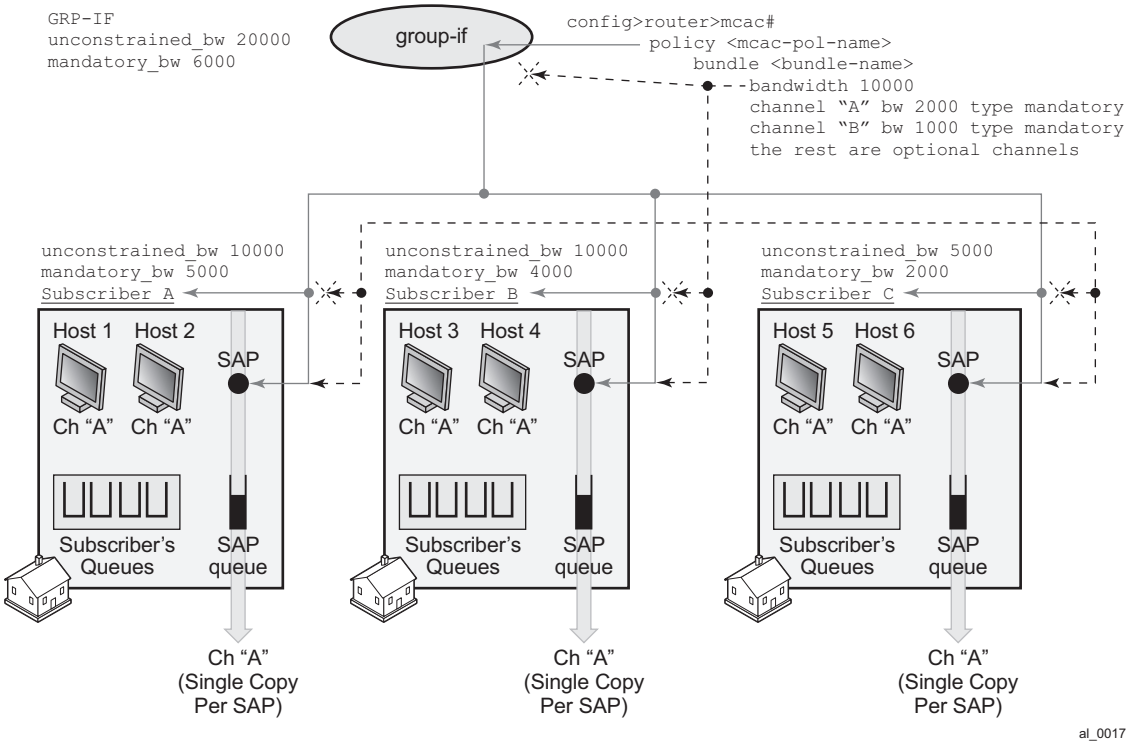

Figure 63 depicts behavior in per-host replication mode. MCAC policy inheritance flow is the same as in the previous example with the difference that the bundle limit has no effect at all. Each host generates its own copy of the same multicast stream that is flowing via subscriber queues and not the SAP queue. Since each of the copies counts towards the subscriber or group-interface bandwidth limits, the multicast bandwidth consumption is higher in this example. This needs to be reflected in the configured multicast bandwidth limits. For example, the group-interface mandatory bandwidth limit is increased to 12 Mb/s.

In our example, 6 subscriber hosts are still watching the same channel but now the number of replications is doubled from previous example. So the final tally for our MCAC bandwidth limit is as follows:

- Subscriber A - 1 Mb/s is still reserved for mandatory channels and 5 Mb/s for optional channels (first come first serve). All this assume that SAP and grp-if bandwidth checks pass.
- Subscriber B - No reserved mandatory bandwidth is left. 6 Mb/s is s till left for optional channels (first come first serve). All this assume that SAP and grp-if bandwidth checks pass.

- Subscriber C - No reserved mandatory bandwidth is left. 1 Mb/s is still left for optional channels (first come first serve). All this assume that SAP and grp-if bandwidth checks pass.
- No reserved bandwidth is left under the group-interface. 8 Mb/s of available multicast bandwidth under the group-interface is still left for optional channels on a first come first serve basis
- Bundle limit on a SAP is irrelevant in this case.

Figure 63 **MCAC Policy Inheritance in Per-HOST Replication Mode**

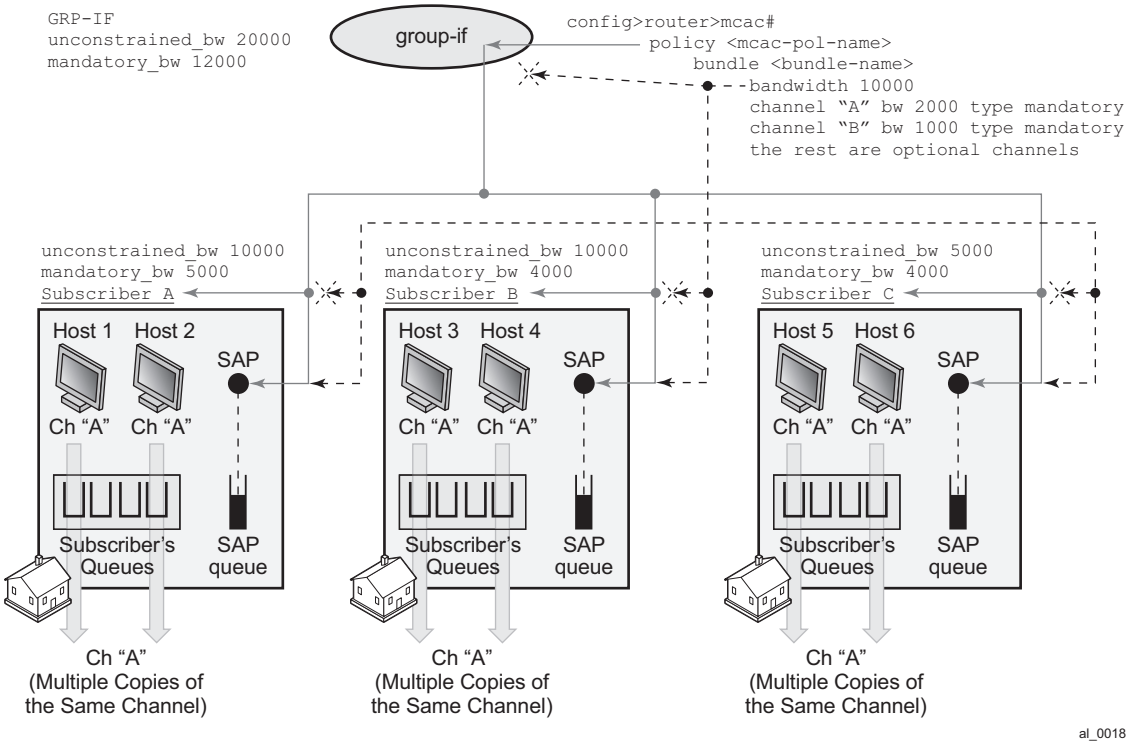

8.4.9 Determining MCAC Policy in Effect

Channel bandwidth definition (via MCAC policy) can be applied under the interface level (group-interface or regular interface):

configure>router/service>igmp/mld>interface/grp-if

The following configuration options can lead to the confusion as to which MCAC policy is in effect:

- The MCAC policy (channel bandwidth definition) can be applied under two different places (grp-if and/or regular intf).
- The same policy is used for (H)MCAC and HQoS Adjust with redirection enabled/disabled.

The general, the rule is that the MCAC policy under the group-interface is always be in effect in cases where redirection is disabled. This is valid for subscriber or group-interface MCAC, hMCAC (subscriber and group-interface) and HQoS Adjust in per SAP replication mode.

If redirection is enabled, the MCAC policy under the group-interface is ignored.

If redirection is enabled, but there is no MCAC policy applied under the redirected interface (redirected interface is the interface to which IGMP/MLD Joins are redirected from subscriber hosts), regardless of whether the MCAC policy under the group-interface is applied or not, then:

- HQoS adjust has no effect.
- MCAC has no effect not only per redirected interface but also per subscriber.

8.4.10 Multicast Filtering

Multicast filtering must be done per session (host) for IPoE and PPPoE. There are two types of filters that are supported:

1. IGMP filters on access ingress. Those filters control the flow of IGMP messages between the host and the BNG. They are applied via the import statement in the igmp-policy. The same filters are used for multicast-redirection policy:

   For IPv4

```
configure
    subscr-mgmt
        igmp-policy <name>
            import <policy-name>
```

   For IPv6

```
configure
    subscr-mgmt
        mld-policy <name>
            import <policy-name>
```

   An example of the filter definition is given below:

```
configure
    router
        policy-options
```

```
begin
prefix-list <pref-name>
    prefix <pref-definition>
policy-statement <name>
    entry 1
        from
            group-address <pref-name>
            source-address <ip>
            protocol igmp
        exit
        action accept
        exit
    exit
    default-action reject
```

2. Regular traffic filters where control multicast traffic flow can be controlled in both directions (ingress/egress). This is supported through ip-filters under the SLA profile.

For IPv6 specifically, up to 15 import policies can be applied to the subscriber IPv6 host. Each of the import policies can be configured as an IPTV package. Depending on the IPTV package the subscriber ordered, a list of import policies is applied at authentication.

At authentication, the first 14 import policies are applied either through LUDB or RADIUS using a list of VSAs. These 14 import policies are not read in any particular order and to make the list deterministic as a black list or a white list, the import policy inside the MLD policy must be configured. The import policy inside the MLD policy is always applied as the last policy which determines whether the list of 14 policies is a white list or a black list. For example, to make the list of 15 import policies a white list, the first 14 policies should have **action** configured as **accept** and the last import policy inside the MLD policy should have default action configured as **reject**. For a black list, the first 14 policies should have **action** configured as **reject** and the last import policy inside the MLD policy should have default action configured as **accept**. It is recommended to configure the MLD import policy with a default action without an import list which is used to determine if the list is a black or white list. If the last import policy is overridden, it changes the MLD policy which in turn overrides the subscriber profile.

It is possible to make a mixed black and white list, where the general rule is that the import policy inside the MLD policy is always the last to be applied to the subscriber.

The list of MLD import policies can be overridden either through RADIUS CoA or through the **tools**>**perform**>**subscriber-management**>**coa** command. There are two VSAs for overriding the list of import policies. The first VSA, Alc-Mld-Import-Policy, is used to completely override the first 14 MLD policies. A RADIUS CoA can contain up to 14 VSAs to override the existing list. Due to limitations of the length of a CLI command, up to five VSAs can be passed to override the existing list when using the **tools** command. The second VSA, Alc-Mld-Import-Policy-Modif, allows the

addition or subtraction of a single MLD import policy. To add an import policy to the existing list, use a prefix of ‘a:’ and to remove an import policy, use a prefix of ‘s:’, for example, ‘a:sport-pack1-hd’ or ‘s:movie-pack1-hd’. Multiple import policies can be added or removed within a single CoA by using multiple VSAs. Up to five import policies can be added or removed using a single CoA. The last import policy inside the MLD policy can be overridden using the subscriber profile.

If a subscriber is not associated with an MLD policy, the import policies applied either at authentication or overridden in mid-session are always stored while the subscriber is online. The import policy does not take effect as long as the subscriber is not associated with an MLD policy. However, as soon as the subscriber is associated with an MLD policy, the list of import policies stored is instantly applied.

8.4.11 Joining the Multicast Tree

The delivery of multicast to the subscribers-interface in a VPRN environment depends on the multicast deployment model (PIM, mBGP). In each model, a subscriber-interface is treated as a regular CE-PE interface that has registered v4 multicast listeners.

8.4.12 Wholesale/Retail Requirements

Multicast support on subscriber interfaces is supported in both wholesale/retail models:

- Wholesale/retail VPRN (IPoE and PPPoE)
- LAC/LNS (PPPoE only)

  This model is shown in Figure 64.

Figure 64 **Wholesale/Retail Multicast Support**

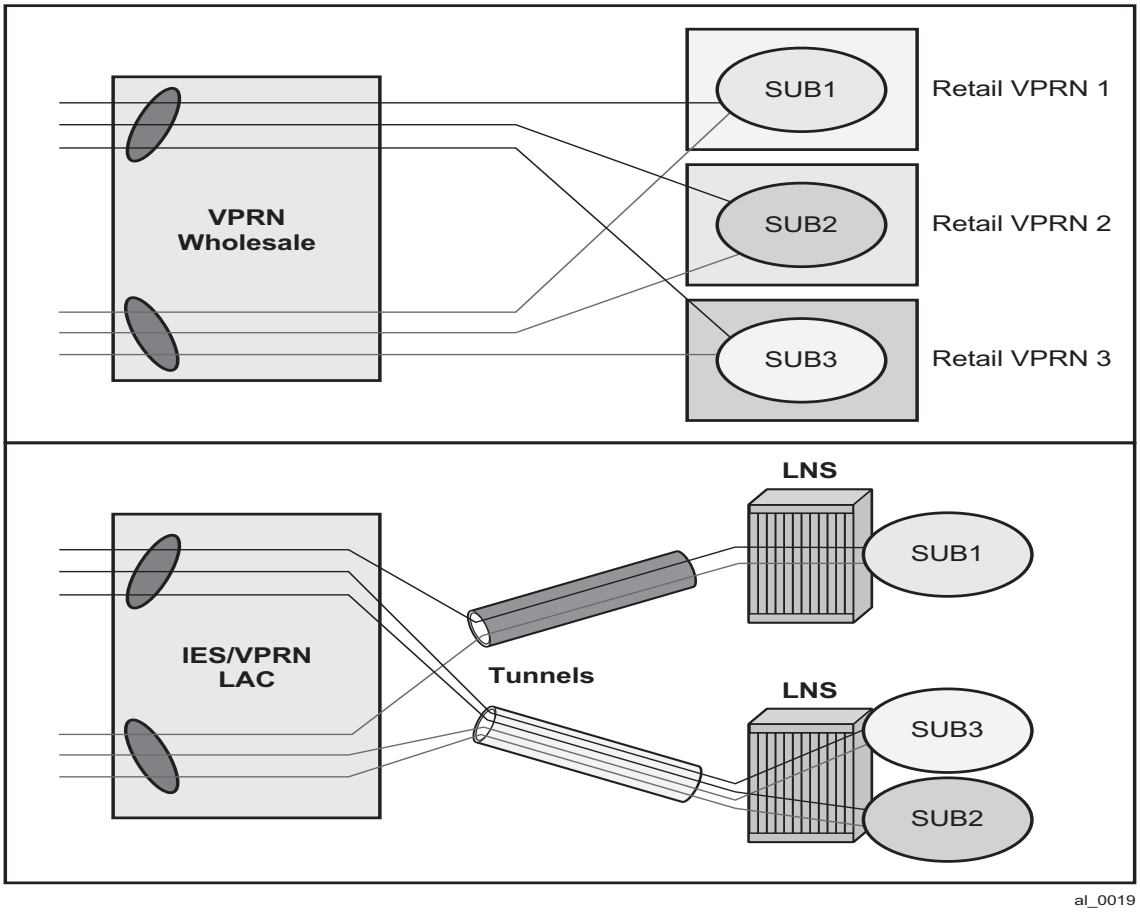

The distinction between these two models is that in the case of LAC/LNS, the replication is performed further up in the network on an LNS node. This means that the traffic between LAC and LNS is multiplied by the amount of replications.

8.4.13 QoS Considerations

In per-sap replication mode (which is applicable only to IPoE subscribers), multicast traffic is forwarded through the SAP queue which is outside of the subscribers queues and therefore not accounted in subscriber aggregate rate limit. HQoS Adjust is used to remedy this situation.

In case that the SAP queue is removed from the static SAP in IPoE 1:1 model (with profiled-only-traffic command), multicast traffic flows via internal queues which cannot be tied into a port-scheduler as part of HQoS. Consequently, the **port-scheduler** *max-rate* as defined in the port scheduler policy is used only to rate limit unicast traffic. In other words, the *max-rate* value in the port scheduler policy must be lowered for the amount of anticipated multicast traffic that flows via the port where the port scheduler policy is applied.

A similar logic applies to the per-sap replication mode on dynamic SAPs (MSAPs) even if the SAP queue is not removed. Although the multicast traffic is flowing via the SAP queue in this case, the SAP qos policy on MSAP cannot be changed from the default one. The default QoS policy on a SAP contains a single queue that is not parented by the port-scheduler.

Those restrictions do not apply to static SAPs where the SAP QoS policy can be customized and its queues consequently tied to the port-scheduler.

8.4.14 Redundancy Considerations

The subscriber can receive multicast content through the subscriber SAP, the redirected interface, or a combination of both.

Multicast redundancy is only supported on a MC-LAG topology. Multicast traffic can be delivered over the subscriber SAP, the redirected interface, or a combination of both.

Subscriber IGMP states can be synchronized across multiple 7450 ESS and 7750 SR nodes in order to ensure minimal interruption of (video delivery) service during network outages. The IGMP/MLD state of a subscriber-host in a 7450 ESS and 7750 SR node is tied to the state of the underlying MC-LAG protection mechanism. For example, IGMP states is activated only for subscribers that are anchored under the group-interfaces with master SRRP state or under an active MC-LAG port.

For multicast redirection on a MC-LAG topology, it must be ensured that the redirected interface (the interface to which multicast forwarding is redirected) is under the same MC-LAG as the subscriber. Otherwise, IGMP states on the redirected interface is derived independently of the IGMP states for the subscriber from which IGMP/MLD messages are redirected.

The IGMP/MLD synchronization process in conjunction with underlying access protection mechanisms works as follows:

- IGMP/MLD states for the subscriber updates only if IGMP/MLD messages (Joins/Leaves/Reports, and so on) are received:

– Directly from the downstream node on an active MC-LAG link. This is valid irrespective of the IGMP querier status for the subscriber.

In all other cases, assuming that some protection mechanism in the access is present, the IGMP/MLD messages are discarded and consequently no IGMP/MLD state is updated. Similar logic applies to regular Layer 3 interfaces, where SRRP is replaced with VRRP.

- Once the subscriber IGMP/MLD state is updated as a result of directly received IGMP/MLD message on an active subscriber (SRRP master of active MC-LAG), the sync IGMP/MLD message is sent to the standby subscriber over the Multi-Chassis Synchronization protocol. Synchronized IGMP states is populated in Multi-chassis Synchronization (MCS) DB in all pairing 7450 ESS and 7750 SR nodes.
- In case that a IGMP/MLD sync (MCS) message is received from the peering node, the IGMP state for the standby subscriber is updated in the MCS DB but it is not downloaded into the forwarding plane unless there is a switchover. In case that the IGMP/MLD sync message is received for the active subscriber, the message is discarded.
- IGMP/MLD queries are sent out only by IGMP/MLD querier. In a MC-LAG environment, it is the node with the active MC-LAG link.

Note: MC-LAG is usually configured with SRRP and the SRRP state is derived from the MC-LAG.

- IGMP/MLD states from the MCS DB are:
  – Activated on non-querier subscriber in case that neither SRRP nor MC-LAG is deployed. It is assumed that the querier subscriber has received the original IGMP/MLD message and consequently sent the IGMP/MLD MCS Sync to the non-querier (standby). Non-querier interface accepts the MCS sync message and also it propagate these IGMP/MLD states to PIM.

  The querier subscriber will not accept the IGMP/MLD update from the MCS database.

  – Aware of the state of MC-LAG. As soon as the standby MC-LAG becomes active, the IGMP/MLD states is activated and they will be propagated to PIM. Traffic is forwarded as soon as multicast streams are delivered to the node and the IGMP states under the subscriber are activated. On a standby MC-LAG, IGMP states are not propagated from the MCS DB to PIM and consequently subscribers.

  - Aware of the SRRP state. Since the subscriber with SRRP Master state is considered active, the states are propagated to PIM as well. On standby SRRP, IGMP states are be propagated from MCS DB to PIM and consequently to subscribers.
- For MC-lag setups, once the switchover is triggered via MC-LAG or SRRP, the IGMP/MLD states from MCS DB on the newly active MC-LAG node or subscriber under the newly SRRP Master is sent to PIM and consequently to the forwarding plane effectively turning on multicast forwarding.

An active and standby subscriber refers to the state of underlying protection mechanism (active MC-LAG).

Note: The subscribers themselves are always instantiated (or active) on both nodes. However, traffic forwarding over those subscribers is driven by the state of the underlying protection mechanism (MC-LAG). Hence the terms active and standby subscriber.

Note: In subscriber environment, SRRP should be always activated in dual-homing scenario. SRRP in subscriber environment ensures that downstream traffic is forwarded via the same node that is forwarding upstream traffic. In this fashion, accounting and QoS for the subscriber are consolidated within a single node.

To summarize, in multi-chassis environment with subscribers, IGMP synchronization enabled and an access layer protection mechanism in place (MC-LAG), the behavior for is the following:

- IGMP/MLD states are synchronized between the chassis.
- On a MC-LAG setup, only the SRRP master or active MC-LAG forwards downstream multicast traffic.
- Length of outage during the switchover is determined by the detection and recovery of the underlying protection mechanism (MC-LAG or MCS) in addition to local propagation of IGMP/MLD states from MCS DB to PIM and consequently to forwarding plane.

Note: IGMP/MLD states can be statically configured on both redundant nodes in order to attract multicast traffic from upstream and therefore minimize outage during the switchover.

8.4.14.1 Redirection Considerations

The redirection policy has two options: to redirect only a certain set of multicast groups to the redirect interface or redirect all multicast to the redirect interface. The redirect policy is source agnostic.

On a MC-LAG setup, for redirection and MCS to work simultaneously in predictable manner, the redirected interface and the corresponding subscribers have to be protected by the same MC-LAG. This binds the redirected interfaces and the subscriber-hosts to the same physical port(s).

The following describe some guidelines for a MC-LAG setup:

- The active subscriber replicates its received IGMP/MLD message to the redirected Layer 3 interface. The Layer 3 redirected interface accepts this message:
  - Independently of the corresponding VRRP state if MC-LAG is not used.
  - Only if the Layer 3 interface is IGMP querier.
  - MC-LAG is used and in active state.
- In all other cases the IGMP message under the Layer 3 redirected interface is be rejected.

Note: Layer 3 redirected interface can also receive IGMP message directly from the downstream node in case that IGMP forking in the access node is activated.

- The Layer 3 redirected interface will not accept the IGMP state update from the MCS DB unless the Layer 3 interface is a non-querier.
- In case that the Layer 3 redirected interface is part of MC-LAG, the IGMP state update sent to it via MCS DB is accepted only during the transitioning phase from standby to active MC-LAG state.

  Briefly, IGMP states on Layer 3 interface are not VRRP aware. However, they are MC-LAG aware.

8.4.15 Query Intervals for Multicast

The query interval commands—**query-interval**, **query-last-listener-interval**, **query-last-member-interval**, and **query-response-interval** are configurable in the following command contexts:

- router and VPRN service IGMP and MLD command contexts

- router and VPRN service IGMP and MLD group interface command contexts
- IGMP policy and MLD policy command contexts

The IGMP policy and MLD policy query intervals are only applicable to ESM host-based queries. Group interface query intervals are only applicable when the **no sub-hosts-only** command is configured. This triggers all group interface SAP(s) to send multicast queries. IGMP policy, MLD policy, and group interface timers use the router IGMP and MLD timers as a default fallback (see ESM Host-based Queries and Group Interface-based Queries).

For IGMP, the **query-interval** must be configured to be longer than the **query-last-member-interval** and **query-response-interval**. For MLD, the **query-interval** must be configured to be longer than the **query-last-listener-interval** and **query-response-interval**.

8.4.15.1 ESM Host-based Queries

An ESM host associated with an IGMP policy or an MLD policy is sent host-based queries. A host-based query uses the intervals defined under the IGMP policy or MLD policy first, and if none are configured, it uses the interval values configured under the router IGMP or MLD context as a fallback. The router IGMP and MLD interval values are set to default values if none are configured. It is possible for the IGMP or MLD policy to only configure some of the query intervals while leaving other query intervals to use the values configured under the router IGMP or MLD context. For example, if the IGMP policy only has the **query-interval** configured to 11, the remaining **query-last-member-interval** and **query-response-interval** configurations uses values 1 and 10, respectively, defined under **router igmp**. It is also possible to configure an invalid combination of query timers when using both the IGMP policy configured intervals and intervals under **router igmp**. Using the example above, if the IGMP policy only sets the **query-interval** to 5, while the **query-response-interval** under **router igmp** is set to 10, this becomes an invalid combination because the response interval is longer than the query interval. When the system detects an invalid combination, it falls back and use all interval values configured under the **router igmp** or **router mld** context and ignores the interval configured under the IGMP or MLD policy. It is highly recommended that all three query intervals be configured together under the IGMP or MLD policy. To see the active values that each host is using for queries, use the **show router igmp hosts detail** or **show router mld host** command.

8.4.15.2 Group Interface-based Queries

When the group interface is configured with **no sub-hosts-only** (applicable only to IPoE multicast), each SAP configured under that group interface sends out SAP-based queries. In this case, a single query is sent out of the SAP. The query intervals uses the values configured under the group interface. If there are no values configured under the group interface, the router IGMP or MLD values are used. When using the **no sub-hosts-only** command, IGMP and MLD policy are not required. To see the active values for the SAP-based queries, use the **show router igmp group-interface detail** or **show router mld group-interface detail** command.

8.5 Configuring Triple Play Multicast Services with CLI

This section provides information to configure multicast parameters in a Triple Play network using the command line interface.

8.5.1 Configuring IGMP Snooping in the BSA

8.5.1.1 Enabling IGMP Snooping in a VPLS Service

8.5.1.2 IGMPv3 Multicast Routers

When multicast routers use IGMPv3, it is sufficient to only enable IGMP snooping, without any further modification of parameters.

The following displays an example of an IGMP snooping configuration:

```
A:ALA-48>config>service>vpls# info
----------------------------------------------
            igmp-snooping
                no shutdown
            exit
            no shutdown
----------------------------------------------
A:ALA-48>config>service>vpls#
```

8.5.1.3 With IGMPv1/2 Multicast Routers

When the multicast routers do not support IGMPv3, some timing parameters must be configured locally in the SR-series routers.

Note: All routers in the multicast network must use the same values for these parameters.

The following displays an example of a modified IGMP snooping configuration:

```
A:ALA-48>config>service>vpls# info
----------------------------------------------
            stp
                shutdown
            exit
            igmp-snooping
                query-interval 60
                robust-count 5
                no shutdown
            exit
            no shutdown
----------------------------------------------
A:ALA-48>config>service>vpls#
```

8.5.1.4 Modifying IGMP Snooping Parameters

For interoperability with some multicast routers, the source IP address of IGMP group reports can be configured. Use the following CLI syntax to customize this IGMP snooping parameter:

The following displays an example of a modified IGMP snooping configuration:

```
A:ALA-48>config>service>vpls# info
----------------------------------------------
            stp
                shutdown
            exit
            igmp-snooping
                query-interval 60
                robust-count 5
                report-src-ip 10.20.20.20
                no shutdown
            exit
            no shutdown
----------------------------------------------
A:ALA-48>config>service>vpls#
```

8.5.1.5 Modifying IGMP Snooping Parameters for a SAP or SDP

Use the following CLI syntax to customize IGMP snooping parameters on an existing SAP. Commands for spoke or mesh SDPs are identical.

CLI Syntax:

```
config>service# vpls service-id
sap sap-id
     igmp-snooping
          fast-leave
          import policy-name
          last-member-query-interval interval
          max-num-groups max-num-groups
```

```
mrouter-port
query-interval interval
query-response-interval interval
robust-count count
send-queries
```

To enable and customize sending of IGMP queries to the hosts:

Example:

```
config>service# vpls 1
config>service>vpls# sap 1/1/3:0
config>service>vpls>sap# igmp-snooping
config>service>vpls>sap>snooping# send-queries
config>service>vpls>sap>snooping# query-interval 100
config>service>vpls>sap>snooping# query-response-
 interval 60
config>service>vpls>sap>snooping# robust-count 5
config>service>vpls>sap>snooping# exit
config>service>vpls>sap# no shutdown
```

To customize the leave delay:

Example:

```
config>service# vpls 1
config>service>vpls# sap 1/1/1:1
config>service>vpls>sap# igmp-snooping
config>service>vpls>sap>snooping# last-member-query-
 interval 10
config>service>vpls>sap>snooping# no fast-leave
config>service>vpls>sap>snooping# exit
config>service>vpls>sap# exit
```

To enable Fast Leave:

Example:

```
config>service# vpls 1
config>service>vpls# sap 1/1/1:1
config>service>vpls>sap# igmp-snooping
config>service>vpls>sap>snooping# no last-member-query-
 interval
config>service>vpls>sap>snooping# fast-leave
config>service>vpls>sap>snooping# exit
config>service>vpls>sap# exit
```

To limit the number of streams that a host can join:

Example:

```
config>service# vpls 1
config>service>vpls# sap 1/1/1:1
config>service>vpls>sap# igmp-snooping
config>service>vpls>sap>snooping# max-num-groups 4
config>service>vpls>sap>snooping# exit
```

```
config>service>vpls>sap# exit
```

To enable sending group reports on a SAP to standby multicast routers:

Example:
```
config>service# vpls 1
config>service>vpls# sap 1/1/1:1
config>service>vpls>sap# igmp-snooping
config>service>vpls>sap>snooping# mrouter-port
config>service>vpls>sap>snooping# exit
config>service>vpls>sap# exit
```

The following example displays the modified IGMP snooping configuration on a SAP:

```
A:ALA-48>config>service>vpls>sap>snooping# info detail
----------------------------------------------
                    no fast-leave
                    no import
                    max-num-groups 4
                    last-member-query-interval 10
                    no mrouter-port
                    query-interval 100
                    query-response-interval 60
                    robust-count 5
                    send-queries
----------------------------------------------
A:ALA-48>config>service>vpls>sap>snooping#
```

8.5.2 Configuring Static Multicast Groups on a SAP or SDP

Use the following CLI syntax to add static group membership entries on an existing SAP (commands for spoke or mesh SDPs are identical):

The following displays an example of a static IGMP snooping configuration on a SAP:

```
A:ALA-48>config>service>vpls>sap# info
----------------------------------------------
                max-nbr-mac-addr 4
                igmp-snooping
                    fast-leave
                    mrouter-port
                    static
                        group 239.0.10.10
                            source 10.10.10.1
                            source 10.10.10.2
                        exit
                    exit
                exit
----------------------------------------------
A:ALA-48>config>service>vpls>sap#
```

8.5.2.1 Enabling IGMP Group Membership Report Filtering

Routing policies can be defined to limit the multicast channels that can be joined by a host. For example, it is possible to define a policy listing a group of multicast streams (for example, 'basic' containing a basic set of TV channels or 'extended' containing a more extended set of TV channels), and to apply this policy to subscribers of IGMP snooping (SAPs or SDPs).

The following displays an example of a configuration to import a routing policy on a SAP:

```
A:ALA-48>config>service>vpls# info
----------------------------------------------
            stp
                shutdown
            exit
            igmp-snooping
                query-interval 60
                robust-count 5
                report-src-ip 10.20.20.20
                no shutdown
            exit
            sap 1/1/3:0 create
                igmp-snooping
                    query-interval 100
                    query-response-interval 60
                    robust-count 5
                    send-queries
                exit
            exit
            sap 1/1/3:22 create
                max-nbr-mac-addr 4
                igmp-snooping
                    fast-leave
                    import "test_policy"
                    mrouter-port
                    static
                        group 239.0.10.10
                            source 10.10.10.1
                            source 10.10.10.2
                        exit
                    exit
                exit
            exit
            no shutdown
----------------------------------------------
A:ALA-48>config>service>vpls#
```

For details configuring a routing policy, refer to the Configuring Route Policies section in the 7450 ESS, 7750 SR, 7950 XRS, and VSR Router Configuration Guide.

The following shows a sample routing policy configuration accepting IGMP messages for only five multicast channels:

```
A:ALA-48>config>router>policy-options# info
----------------------------------------------
            prefix-list "basic_channels"
                prefix 239.10.0.1/32 exact
                prefix 239.10.0.2/32 exact
                prefix 239.10.0.3/32 exact
                prefix 239.10.0.4/32 exact
                prefix 239.10.0.5/32 exact
            exit
            policy-statement "test_policy"
                description "basic set of 5 multicast channels"
                entry 1
                    from
                        group-address "basic_channels"
                    exit
                    action accept
                    exit
                exit
                default-action reject
            exit
----------------------------------------------
A:ALA-48>config>router>policy-options#
```

8.5.2.1.1 Enabling IGMP Traffic Filtering

For security, it might be advisable to only allow multicast traffic into the SR-series routers from recognized multicast routers and servers. Multicast packets arriving on other interfaces (for example, customer-facing SAPs or spoke SDPs) can be filtered out by defining an appropriate IP filter policy.

For details on how to configure a filter policy, see section Creating an IP Filter Policy in the Router Configuration Guide.

The following example shows a sample IP filter policy configuration dropping all multicast traffic:

```
A:ALA-48>config>filter>ip-filter# info
----------------------------------------------
        ip-filter 1 create
            entry 1 create
                match
                    dst-ip 239.0.0.0/24
                exit
                action accept
            exit
            entry 2 create
                match
                    dst-ip 239.0.0.0/4
                exit
                action drop
            exit
        exit
----------------------------------------------
A:ALA-48>config>filter>ip-filter#
```

The following example shows how to apply this sample IP filter policy to a SAP:

```
A:ALA-48>config>service>vpls # info
----------------------------------------------
         sap 1/1/1:1
             ingress
                 filter ip 1
                 exit
             exit
         exit
----------------------------------------------
A:ALA-48>config>service>vpls>snooping#
```

8.5.3 Configuring Multicast VPLS Registration (MVR)

Use the following CLI syntax to configure Multicast VPLS Registration. The first step is to register a VPLS as a multicast VPLS.

CLI Syntax:

```
config>service# vpls service-id
igmp-snooping
     mvr
          no shutdown
          description description
          group-policy policy-name
```

Example:

```
config>service# vpls 1000
config>service>vpls# igmp-snooping
config>service>vpls>snooping# mvr
config>service>vpls>snooping>mvr# no shutdown
config>service>vpls>snooping>mvr# description "MVR VPLS"
config>service>vpls>snooping>mvr# group-policy
 "basic_channels_policy"
```

The second step is to configure a SAP to take the multicast channels from the registered multicast VPLS.

CLI Syntax:

```
config>service# vpls service-id
sap sap-id
     igmp-snooping
          mvr
               from-vpls vpls-id
```

Example:

```
config>service# vpls 1
config>service>vpls# sap 1/1/1:100
config>service>vpls>sap# igmp-snooping
config>service>vpls>snooping# mvr
config>service>vpls>snooping>mvr# from-vpls 1000
```

For MVR by proxy also the destination SAP for the multicast channels should be configured.

CLI Syntax:
```
config>service# vpls service-id
sap sap-id
     igmp-snooping
          mvr
               from-vpls vpls-id
               to-sap sap-id
```

Example:
```
config>service# vpls 1
config>service>vpls# sap 1/1/1:100
config>service>vpls>sap# igmp-snooping
config>service>vpls>snooping# mvr
config>service>vpls>snooping>mvr# from-vpls 1000
config>service>vpls>snooping>mvr# to-sap 1/1/1:200
```

8.5.4 Configuring IGMP, MLD, and PIM in the BSR

Note: Configuring IGMP, MLD, and PIM in the BSR is supported by the 7750 SR only.

Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Multicast Routing Protocols Guide* for information about multicast and the commands required to configure basic IGMP and PIM parameters.

8.5.4.1 Enabling IGMP

The following displays an example of enabled IGMP.

```
A:LAX>>config>router/service>vprn info detail
...
#------------------------------------------
echo "IGMP Configuration"
#------------------------------------------
        igmp
            query-interval 125
            query-last-member-interval 1
            query-response-interval 10
            robust-count 2
            no shutdown
        exit
#------------------------------------------
```

```
...
A:LAX>>config>system#
```

8.5.4.2 Configuring IGMP Interface Parameters

The following example displays an IGMP configuration:

```
A:LAX>config>router/service>vprn>igmp# info
----------------------------------------------
        interface "lax-sjc"
        exit
        interface "lax-vls"
        exit
        group-interface "Mcast-subscribers"
        exit
----------------------------------------------
A:LAX>config>router>igmp# exit
```

8.5.4.3 Configuring Static Parameters

The following example displays a configuration to add IGMP a static multicast source:

```
A:LAX>config>router/service>vprn>igmp# info
----------------------------------------------
        interface "lax-sjc"
        exit
        interface "lax-vls"
            static
                group 239.255.0.2
                    source 172.22.184.197
                exit
            exit
        exit
        group-interface "mcast-subscribers"
        interface "lax-vls"
        static
        group 239.255.0.2
        source 172.22.184.197
        exit
        exit
        exit
        exit----------------------------------------------
A:LAX>config>router>igmp#
```

:The following example displays the configuration to add a IGMP static starg entry:

```
A:LAX>config>router/service>vprn>igmp# info
```

```
-----------------------------------------------
        interface "lax-sjc"
            static
                group 239.1.1.1
                    starg
                exit
            exit
        exit
        interface "lax-vls"
            static
                group 239.255.0.2
                    source 172.22.184.197
                exit
            exit
        exit
        group-interface "mcast-subscribers"
        static
        group 239.1.1.1
        starg
        exit
        exit
        exit
-----------------------------------------------
A:LAX>config>router>igmp#
```

8.5.4.4 Configuring SSM Translation

The following displays an SSM translation configuration:

```
A:LAX>config>router/service>vprn>igmp# info
-----------------------------------------------
        ssm-translate
            grp-range 239.255.0.1 231.2.2.2
                source 10.1.1.1
            exit
        exit
        interface "lax-sjc"
            static
                group 239.1.1.1
                    starg
                exit
            exit
        exit
        interface "lax-vls"
            static
                group 239.255.0.2
                    source 172.22.184.197
                exit
            exit
        exit
        group-interface "mcast-subscribers"
        static
        group 239.1.1.1
        starg
        exit
```

```
        exit
        exit
-----------------------------------------------
A:LAX>config>router/service>vprn>igmp# exit
```

8.5.4.5 Enabling MLD

The following displays an example of enabled MLD.

```
A:LAX>>config>router/service>vprn # info detail

#-----------------------------------------
echo "MLD Configuration"
#-----------------------------------------
mld
query-interval 125
query-last-member-interval 1
query-response-interval 10
robust-count 2
no shutdown
exit
#-----------------------------------------
...
A:LAX>>config>router/service>vprn #
```

8.5.4.6 Configuring MLD Interface Parameters

The following example displays an MLD configuration:

```
A:LAX>config>router/service>vprn>mld# info
-----------------------------------------------
interface "lax-sjc"
exit
interface "lax-vls"
exit
group-interface "mcast-subscribers"
exit
-----------------------------------------------
A:LAX>config>router/service>vprn>mld#exit
```

8.5.4.7 Configuring Static Parameters

The following example displays a configuration to add MLD a static multicast source:

```
A:LAX>config>router/service>vprn>mld# info
-----------------------------------------------
```

```
interface "lax-sjc"
exit
interface "lax-vls"
       static
               group ffe8::1
                      source 2001::1
               exit
       exit
exit
group-interface "mcast-subscribers"
       static
               group ffe8::1
                      source 2001::1
               exit
       exit
exit
-------------------------------------------
A:LAX>config>router/service>vprn>mld#
```

The following example displays the configuration to add a MLD static starg entry:

```
A:LAX>config>router/service>vprn>mld# info
-------------------------------------------
interface "lax-sjc"
       static
               group ffe8::2
                      starg
                      exit
               exit
       exit
interface "lax-vls"
       static
               group ffe8::1
                      source 2001::1
                      exit
               exit
       exit
group-interface "mcast-subscribers"
       static
               group ffe8::2
                      starg
                      exit
               exit
       exit
exit
-------------------------------------------
A:LAX>config>router/service>vprn>mld#
```

8.5.4.8 Configuring SSM Translation

The following displays an SSM translation configuration:

```
A:LAX>config>router/service>vprn>mld# info
```

```
-----------------------------------------------
ssm-translate
        grp-range ff31::1  ff32::1
                source 2001::2
                exit
        exit
interface "lax-sjc"
        static
                group ffe8::2
                       starg
                       exit
                exit
        exit
interface "lax-vls"
        static
                group ffe8::1
                       source 2001::1
                       exit
                exit
        exit
group-interface "mcast-subscribers"
        static
                group ff31::20
                       starg
                       exit
                exit
        exit
exit
-----------------------------------------------
A:LAX>config>router/service>vprn>mld# exit
```

8.5.4.9 Configuring PIM

8.5.4.9.1 Enabling PIM

When configuring PIM, make sure to enable PIM on all interfaces for the routing instance, otherwise multicast routing errors can occur.

The following example displays detailed output when PIM is enabled.

```
A:LAX>>config>router# info detail
...
#------------------------------------------
echo "PIM Configuration"
#------------------------------------------
        pim
            no import join-policy
            no import register-policy
            apply-to none
            rp
                no bootstrap-import
```

```
                no bootstrap-export
                static
                exit
                bsr-candidate
                    shutdown
                    priority 0
                    hash-mask-len 30
                    no address
                exit
                rp-candidate
                    shutdown
                    no address
                    holdtime 150
                    priority 192
                exit
            exit
            no shutdown
        exit
#-----------------------------------------
...
A:LAX>>config>system#
```

8.5.4.9.2 Configuring PIM Interface Parameters

The following displays a PIM interface configuration:

```
A:LAX>config>router>pim# info
----------------------------------------------
            interface "system"
            exit
            interface "lax-vls"
            exit
            interface "lax-sjc"
            exit
            interface "p1-ix"
            exit
            rp
                static
                    address 10.22.187.237
                        group-prefix 239.24.24.24/32
                    exit
                    address 10.10.10.10
                    exit
                exit
                bsr-candidate
                    shutdown
                exit
                rp-candidate
                    shutdown
                exit
            exit
----------------------------------------------
A:LAX>config>router>pim#


A:SJC>config>router>pim# info
```

```
-----------------------------------------------
            interface "system"
            exit
            interface "sjc-lax"
            exit
            interface "sjc-nyc"
            exit
            interface "sjc-sfo"
            exit
            rp
                static
                    address 10.22.187.237
                        group-prefix 239.24.24.24/32
                    exit
                exit
                bsr-candidate
                    shutdown
                exit
                rp-candidate
                    shutdown
                exit
            exit
-----------------------------------------------
A:SJC>config>router>pim#


A:MV>config>router>pim# info
-----------------------------------------------
            interface "system"
            exit
            interface "mv-sfo"
            exit
            interface "mv-vlc"
            exit
            interface "p3-ix"
            exit
            rp
                static
                    address 210.22.187.237
                        group-prefix 239.24.24.24/32
                    exit
                exit
                bsr-candidate
                    address 10.22.187.236
                    no shutdown
                exit
                rp-candidate
                    address 10.22.187.236
                    no shutdown
                exit
            exit
-----------------------------------------------
A:MV>config>router>pim#


A:SFO>config>router>pim# info
-----------------------------------------------
            interface "system"
            exit
```

```
            interface "sfo-sjc"
            exit
            interface "sfo-was"
            exit
            interface "sfo-mv"
            exit
            rp
                static
                    address 10.22.187.237
                        group-prefix 239.24.24.24/32
                    exit
                exit
                bsr-candidate
                    address 10.22.187.239
                    no shutdown
                exit
                rp-candidate
                    address 10.22.187.239
                    no shutdown
                exit
            exit
-------------------------------------------
A:SFO>config>router>pim#


A:WAS>config>router>pim# info
-------------------------------------------
            interface "system"
            exit
            interface "was-sfo"
            exit
            interface "was-vlc"
            exit
            interface "p4-ix"
            exit
            rp
                static
                    address 10.22.187.237
                        group-prefix 239.24.24.24/32
                    exit
                exit
                bsr-candidate
                    address 10.22.187.240
                    no shutdown
                exit
                rp-candidate
                    address 10.22.187.240
                    no shutdown
                exit
            exit
-------------------------------------------
A:WAS>config>router>pim#
```

8.5.4.9.3 Importing PIM Join/Register Policies

The import command provides a mechanism to control the (*,g) and (sag) state that gets created on a router. Import policies are defined in the **config>router>policy-options** context. See Configuring PIM Join/Register Policies.

Note: In the import policy, if an action is not specified in the entry then the default-action takes precedence. If no entry matches then the default-action also takes precedence. If no default-action is specified, then the **default-action** is executed.

The following example displays the command usage to apply the policy statement will not allow join messages for group 239.50.50.208/32 and source 192.168.0.0/16 but allows join messages for 192.168.0.0/16, 239.50.50.208:

Example:
```
config>router# pim
config>router>pim# import join-policy "foo"
config>router>pim# no shutdown
```

The following example displays the PIM configuration:

```
A:LAX>config>router>pim# info
----------------------------------------------
            import join-policy "foo"
            interface "system"
            exit
            interface "lax-vls"
            exit
            interface "lax-sjc"
            exit
            interface "p1-ix"
            exit
            rp
                static
                    address 10.22.187.237
                        group-prefix 239.24.24.24/3
                    exit
                    address 10.10.10.10
                    exit
                exit
                bsr-candidate
                    shutdown
                exit
                rp-candidate
                    shutdown
                exit
            exit
----------------------------------------------
A:LAX>config>router>pim#
```

8.5.4.9.4 Configuring PIM Join/Register Policies

Join policies are used in Protocol Independent Multicast (PIM) configurations to prevent the transportation of multicast traffic across a network and the dropping of packets at a scope at the edge of the network. PIM Join filters reduce the potential for denial of service (DoS) attacks and PIM state explosion—large numbers of Joins forwarded to each router on the RPT, resulting in memory consumption.

*,g or s,g is the information used to forward unicast or multicast packets.

- **group-address** matches the group in join/prune messages
  group-address 239.55.150.208/32 exact
- **source-address** matches the source in join/prune messages
  source-address 192.168.0.0/16 longer
- **interface** matches any join message received on the specified interface
  interface port 1/1/1
- **neighbor** matches any join message received from the specified neighbor
  neighbor 1.1.1.1

The following configuration example will not allow join messages for group 239.50.50.208/32 and source 192.168.0.0/16 but allows join messages for 192.168.0.0/16, 239.50.50.208.

```
A:ALA-B>config>router>policy-options# info
----------------------------------------------
...
            policy-statement "foo"
                entry 10
                    from
                        group-address "239.50.50.208/32"
                        source-address 192.168.0.0
                    exit
                    action reject
                exit
            exit
            policy-statement "reg-pol"
                entry 10
                    from
                        group-address "239.0.0.0/8"
                    exit
                    action accept
                    exit
                exit
            exit
...
----------------------------------------------
A:ALA-B>config>router>policy-options#
```

8.5.4.10 Configuring Bootstrap Message Import and Export Policies

Bootstrap import and export policies are used to control the flow of bootstrap messages to and from the RP.

The following configuration example specifies that no BSR messages received or sent out of interface port 1/1/1.

```
:A:ALA-B>config>router>policy-options# policy-statement pim-import
:A:ALA-B>config>router>policy-options>policy-statement$ entry 10
:A:ALA-B>config>router>policy-options>policy-statement>entry$ from
:A:ALA-B>config>router>policy-options>policy-statement>entry>from$ interface port1/
1/1/
:A:ALA-B>config>router>policy-options>policy-statement>entry>from$ exit
:A:ALA-B>config>router>policy-options>policy-statement>entry# action reject
:A:ALA-B>config>router>policy-options>policy-statement>entry# exit
:A:ALA-B>config>router>policy-options>policy-statement# exit

:A:ALA-B>config>router>policy-options# policy-statement pim-export
:A:ALA-B>config>router>policy-options>policy-statement$ entry 10
:A:ALA-B>config>router>policy-options>policy-statement>entry$ to
:A:ALA-B>config>router>policy-options>policy-statement>entry>to$
```

8.6 Triple Play Multicast Command Reference

8.6.1 Command Hierarchies

- Multicast Management Commands
- Multicast Info Policy Bundle Commands
- Multicast Service Commands
- Ingress Multicast Path Management Commands
- Multicast Query Timer Commands
- Show Commands
- Debug Commands

8.6.1.1 Multicast Management Commands

```
config
    — mcast-management
        — bandwidth-policy policy-name [create]
        — no bandwidth-policy policy-name
            — admin-bw-threshold kilo-bits-per-second
            — no admin-bw-threshold
            — description description
            — no description
            — falling-percent-reset percent-of-highest
            — no falling-percent-reset
            — mcast-pool percent-of-total percent-of-buffers resv-cbs percent-of-pool
                slope-policy policy-name
            — no mcast-pool
        — chassis-level
            — [no] mmrp-impm-override
            — [no] per-mcast-plane-capacity
                — mcast-capacity primary-percentage secondary secondary-
                    percentage
                — no mcast-capacity
                — redundant-mcast-capacity primary-percentage secondary
                    secondary-percentage
                — no redundant-mcast-capacity
                — total-capacity capacity
                — no total-capacity
            — [no] round-robin-inactive-records
        — mcast-reporting-dest mcast-reporting-dest-name [create]
        — no mcast-reporting-dest mcast-reporting-dest-name
            — description description
            — no description
```

```
— dest-ip-addr ip-address
— no dest-ip-addr
— max-tx-delay delay
— no max-tx-delay
— udp-dst-port port
— no udp-dst-port
— [no] shutdown
```

8.6.1.2 Multicast Info Policy Bundle Commands

```
config
    — mcast-management
        — multicast-info-policy policy-name [create]
        — no multicast-info-policy policy-name
            — bundle bundle-name [create]
            — no bundle bundle-name
                — admin-bw kbps
                — no admin-bw
                — bw-activity {use-admin-bw | dynamic [falling-delay seconds]} [black-
                    hole-rate kbps]
                — no bw-activity
                — channel ip-address [ip-address] [create]
                — no channel ip-address [ip-address]
                    — admin-bw kbps
                    — no admin-bw
                    — bw-activity {use-admin-bw | dynamic [falling-delay seconds]}
                        [black-hole-rate kbps]
                    — no bw-activity
                    — explicit-sf-path {primary | secondary}
                    — no explicit-sf-path
                    — keepalive-override keepalive-timer
                    — no keepalive-override
                    — preference preference-level
                    — no preference
                    — primary-tunnel-interface {rsvp-p2mp | ldp-p2mp p2mp-id}
                        lsp-name sender ip-address
                    — no primary-tunnel-interface
                    — source-override ip-address [create]
                    — no source-override ip-address
                        — admin-bw kbps
                        — no admin-bw
                        — bw-activity {use-admin-bw | dynamic [falling-delay
                            seconds]} [black-hole-rate kbps]
                        — no bw-activity
                        — explicit-sf-path {primary | secondary}
                        — no explicit-sf-path
                        — keepalive-override keepalive-timer
                        — no keepalive-override
                        — preference preference-level
                        — no preference
```

```
                        — primary-tunnel-interface {rsvp-p2mp | ldp-p2mp
                            p2mp-id} lsp-name sender ip-address
                        — no primary-tunnel-interface
                    — cong-priority-threshold preference-level
                    — no cong-priority-threshold
                    — description description
                    — no description
                    — ecmp-opt-threshold preference-level
                    — no ecmp-opt-threshold
                    — explicit-sf-path {primary | secondary}
                    — no explicit-sf-path
                    — keepalive-override keepalive-timer
                    — no keepalive-override
                    — preference preference-level
                    — no preference
                    — primary-tunnel-interface {rsvp-p2mp | ldp-p2mp p2mp-id} lsp-name
                        sender ip-address
                    — no primary-tunnel-interface
                — description description-string
                — no description
```

Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Interface Configuration Guide* for more usage, syntax and command descriptions.

```
config
    — card
        — fp [fp-number]
            — ingress
                — mcast-path-management
                    — bandwidth-policy policy-name
                    — no bandwidth-policy
                    — [no] shutdown
```

8.6.1.3 Multicast Service Commands

```
config
    — service
        — [no] vpls service-id [customer customer-id]  [vpn vpn-id] [m-vpls]
            — igmp-snooping
                — mvr
                    — description description-string
                    — no description
                    — group-policy policy-name
                    — no group-policy
                    — [no] shutdown
                — query-interval seconds
                — no query-interval
                — query-src-ip ipv6-address
                — no query-src-ip
                — report-src-ip ipv6-address
```

```
            — no report-src-ip
            — robust-count robust-count
            — no robust-count
            — [no] shutdown
      — fdb-table-high-wmark high-water-mark
      — no fdb-table-high-wmark
      — fdb-table-low-wmark low-water-mark
      — no fdb-table-low-wmark
      — fdb-table-size table-size
      — no fdb-table-size
      — mld-snooping
            — mvr
                  — description description-string
                  — no description
                  — group-policy policy-name
                  — no group-policy
                  — [no] shutdown
            — query-interval seconds
            — no query-interval
            — query-src-ip ipv6-address
            — no query-src-ip
            — report-src-ip ipv6-address
            — no report-src-ip
            — robust-count robust-count
            — no robust-count
            — [no] shutdown
      — multicast-info-policy policy-name
      — no multicast-info-policy
      — sap sap-id [split-horizon-group group-name]
      — no sap sap-id
            — igmp-snooping
                  — [no] fast-leave
                  — import policy-name
                  — no import
                  — last-member-query-interval interval
                  — no last-member-query-interval
                  — max-num-groups max-num-groups
                  — no max-num-groups
                  — [no] mrouter-port
                  — mvr
                        — from-vpls vpls-id
                        — no from-vpls
                        — to-sap sap-id
                        — no to-sap
                  — query-interval interval
                  — no query-interval
                  — query-response-interval interval
                  — no query-response-interval
                  — robust-count count
                  — no robust-count
                  — [no] send-queries
                  — static
                        — [no] group group-address
                              — [no] source ip-addr
```

```
                    — [no] starg
— mesh-sdp sdp-id[:vc-id] [vc-type {ether | vlan | vpls}]
— no mesh-sdp sdp-id[:vc-id]
    — igmp-snooping
        — [no] fast-leave
        — import policy-name
        — no import
        — last-member-query-interval interval
        — no last-member-query-interval
        — max-num-groups max-num-groups
        — no max-num-groups
        — mcac
            — policy policy-name
            — no policy
            — unconstrained-bw bandwidth mandatory-bw
                mandatory-bw
            — no unconstrained-bw
        — query-interval interval
        — no query-interval
        — query-response-interval interval
        — no query-response-interval
        — robust-count count
        — no robust-count
        — [no] send-queries
        — static
            — [no] group group-address
                — [no] source ip-addr
                — [no] starg
    — mld-snooping
        — [no] disable-router-alert-check
        — [no] fast-leave
        — import policy-name
        — no import
        — last-member-query-interval interval
        — no last-member-query-interval
        — max-num-groups max-num-groups
        — no max-num-groups
        — mvr
            — fast-leave service-id
            — no fast-leave
            — to-sap sap-id
            — no to-sap
        — query-interval seconds
        — no query-interval
        — query-response-interval seconds
        — no query-response-interval
        — robust-count robust-count
        — no robust-count
        — [no] send-queries
— spoke-sdp sdp-id[:vc-id] [vc-type {ether | vlan | vpls}] [split-horizon-group
    group-name]
— no spoke-sdp sdp-id[:vc-id]
    — igmp-snooping
        — [no] fast-leave
```

```
— import policy-name
— no import
— last-member-query-interval interval
— no last-member-query-interval
— max-num-groups max-num-groups
— no max-num-groups
— mvr
    — fast-leave service-id
    — no fast-leave
    — to-sap sap-id
    — no to-sap
— query-interval interval
— no query-interval
— query-response-interval interval
— no query-response-interval
— robust-count count
— no robust-count
— [no] send-queries
— static
    — [no] group group-address
        — [no] source ip-addr
        — [no] starg
— query-interval seconds
— no query-interval
— query-response-interval seconds
— no query-response-interval
— robust-count robust-count
— no robust-count
— [no] send-queries
```

8.6.1.4 Ingress Multicast Path Management Commands

```
config
    — mcast-management
        — bandwidth-policy policy-name [create]
        — no bandwidth-policy policy-name
            — admin-bw-threshold kilo-bits-per-second
            — no admin-bw-threshold
            — description description-string
            — no description
            — falling-percent-reset percent-of-highest
            — no falling-percent-reset
            — mcast-pool percent-of-total percent-of-buffers resv-cbs percent-of-pool
                slope-policy policy-name
            — no mcast-pool
            — t2-paths
                — primary-path
                    — queue-parameters
                        — cbs percentage
                        — no cbs
                        — drop-tail
```

```
                                        — low
                                            — percent-reduction-from-mbs percent-or-
                                                default
                                            — no percent-reduction-from-mbs
                                    — mbs percentage
                                    — no mbs
                            — secondary-path
                                — number-paths number-of-paths [dual-sfm number-of-paths]
                                — queue-parameters
                                    — cbs percentage
                                    — no cbs
                                    — drop-tail
                                        — low
                                            — percent-reduction-from-mbs percent-or-
                                                default
                                            — no percent-reduction-from-mbs
                                    — mbs percentage
                                    — no mbs

config
    — router
        — multicast-info-policy policy-name
        — no multicast-info-policy
        — pim
            — [no] mc-ecmp-balance
            — mc-ecmp-balance-hold minute
            — no mc-ecmp-balance-hold

tools
    — perform
        — router
            — pim
                — mc-ecmp-rebalance [ecmp-opt-threshold]
```

8.6.1.5 Multicast Redirection

```
config
    — router/service vprn
        — policy-options
            — policy-statement
                — entry
                    — action {accept | next-entry | next-policy | drop | reject}
                    — no action
        — igmp
            — interface port:tags
            — mld
                — interface port:tags

config
    — service
        — ies service-id
```

— **vprn** *service-id*
— **interface** *port:tags*

8.6.1.6 Multicast Query Timer Commands

config
— **router**
— **igmp**
— **group-interface**
— [**no**] **disable-router-alert-check**
— **query-interval** *seconds*
— **no query-interval**
— **query-last-member-interval** *seconds*
— **no query-last-member-interval**
— **query-response-interval** *seconds*
— **no query-response-interval**

config
— **router**
— **mld**
— **group-interface**
— [**no**] **disable-router-alert-check**
— **query-interval** *seconds*
— **no query-interval**
— **query-last-listener-interval** *seconds*
— **no query-last-listener-interval**
— **query-response-interval** *seconds*
— **no query-response-interval**

config
— **service**
— **vprn**
— **igmp**
— **group-interface**
— [**no**] **disable-router-alert-check**
— **query-interval** *seconds*
— **no query-interval**
— **query-last-member-interval** *seconds*
— **no query-last-member-interval**
— **query-response-interval** *seconds*
— **no query-response-interval**

config
— **service**
— **vprn**
— **mld**
— **group-interface**
— [**no**] **disable-router-alert-check**
— **query-interval** *seconds*
— **no query-interval**
— **query-last-listener-interval** *seconds*

— **no query-last-listener-interval**
— **query-response-interval** *seconds*
— **no query-response-interval**

8.6.1.7 Show Commands

show
— **mcast-management**
— **bandwidth-policy** [*policy-name*] [**detail**]
— **channel** [**router** *router-instance* | **vpls** *service-id*] [**mda** *slot*[*/mda*]] [**group** *ip-address* [**source** *ip-address*]] [**path** *path-type*] [**detail**]
— **fp** [*slot*-number[*/fp-number*]] [**path** *path-type*]
— **mcast-reporting-dest** [*mcast-reporting-dest-name*]

show
— **service**
— **id** *service-id*
— **host-tracking**
— **saps** [**detail** | **statistics** | **summary** | **mcast-reporting-statistics**]
— **saps sap** *sap-id* [**host** *ip-address*] [**detail** | **statistics** | **summary** | **mcast-reporting-statistics**]

8.6.1.8 Debug Commands

debug
— **mcast-management**
— [**no**] **mcast-reporting-dest** [*dest-name*]
— [**no**] **igmp** [**host** *ip-address*] [**group** *grp-address*]

8.6.2 Command Descriptions

8.6.2.1 Multicast Management Configuration Commands

8.6.2.1.1 Generic Commands

description

Syntax **description** *description-string*

no description

Context config>mcast-mgmt>bandwidth-policy
config>mcast-mgmt>multicast-info-policy
config>mcast-mgmt>multicast-info-policy>bundle
config>mcast-mgmt>mcast-rprt-dest
config>mcast-mgmt>mcast-info-plcy>bundle
config>mcast-mgmt>mcast-info-plcy
config>card>fp>ingress>mcast-mgmt

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes any description string from the context.

Parameters *description-string* — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

shutdown

Syntax [**no**] **shutdown**

Context config>mcast-mgmt>mcast-rprt-dest

Description The shutdown command is used to disable ingress multicast path management for a forwarding plane. The default state is shutdown. When the **no shutdown** command is executed, ingress multicast path management is enabled and the system evaluates all multicast channels on the forwarding plane based on the default or explicit bandwidth policy applied and multicast information policies associated with the channels.

The **no** form of the command is used to enable ingress multicast path management on the forwarding plane.

8.6.2.1.2 Multicast Management Commands

mcast-management

Syntax **mcast-management**

Context config

Description This command enables the context to configure multicast management parameters. The mcast-management CLI node contains the bandwidth-policy and multicast-info-policy definitions. The bandwidth-policy is used to manage the ingress multicast paths into the switch fabric. The multicast-info-policy defines how each multicast channel is handled by the system. The policy may be used by the ingress multicast bandwidth manager, the ECMP path manager and the egress multicast CAC manager.

The mcast-management node always exists and contains the default bandwidth-policy and the default multicast-info-policy. Enter the mcast-management node when editing, deleting or creating a bandwidth policy or multicast info policy. The default bandwidth policy and multicast-info-policy cannot be edited or deleted.

A chassis-level node within multicast-management is used to control the switch fabric multicast planes replication limits. The switch fabric multicast planes are the individual multicast spatial replication contexts available in the system.

mcast-reporting-dest

Syntax **mcast-reporting-dest** *dest-name* [**create**]
no mcast-reporting-dest *dest-name*

Context config>mcast-management

Description This command creates a multicast reporting destination hierarchy in CLI under which parameters defining this destination can be specified. The destination refers to an external node that collects and analyze IGMP events.

The multicast reporting destination is associated with a name that each subscriber can reference in order to send the IGMP related events.

It can be also referenced in the host tracking policy in case that IGMP events are related to the host tracking feature.

The **no** form of the command removes the destination name from the configuration.

Parameters *dest-name* — Specifies the multicast reporting destination name.

dest-ip-addr

Syntax **dest-ip-addr** *ip-addr***ess**
no dest-ip-addr

Context config>mcast-mgmt>mcast-rprt-dest

Description This command specifies the IP address of the external node to which IGMP events are exported. The destination IP address can only be reachable from the global routing table (no vrf access).

The **no** form of the command removes the destination address from the configuration.

Parameters *ip-addr* — Specifies the IP address of the multicast reporting destination.

max-tx-delay

Syntax **max-tx-delay** *deci-seconds*
no max-tx-delay

Context config>mcast-mgmt>mcast-rprt-dest

Description This command specifies the time interval before the packet starts transmitting towards the destination. When an IGMP event is encoded and ready to be transported, a buffer for the packet is allocated (if not already existent). The events are written into this buffer. Along with the initial buffer creation, a timer is started. The trigger for the transmission of the packet is either the TX buffer being filled up to 1400 B, or the timer expiry, whichever comes first.

The **no** form of the command reverts to the default.

Parameters *deci-seconds* — Specifies the maximum delay after which any cached reports are flushed to the reporting destination.

Values 0 to 100

udp-dst-port

Syntax **udp-dst-port** *port*
no udp-dst-port

Context config>mcast-mgmt>mcast-rprt-dest

Description This command specifies the UDP destination port of the external node to which IGMP events are exported.

The **no** form of the command reverts to the default.

Parameters *port* — Specifies the UDP port to send multicast reports.

Values 1 to 65535

8.6.2.1.3 Bandwidth Policy Commands

bandwidth-policy

Syntax **bandwidth-policy** *policy-name* [**create**]

no bandwidth-policy *policy-name*

Context config>mcast-mgmt

Description This command creates a multicast bandwidth policy. Bandwidth policies are used to manage the ingress multicast path bandwidth. Each forwarding plane supports multicast forwarding paths into the switch fabric. By default, two paths are available; the multicast high priority path and the multicast low priority path. Multicast packets are forwarded on either path based on the expedited or non-expedited (best-effort) nature of the queue the packets are scheduled from. The ingress forwarding plane uses the classification rules to determine the forwarding class of each multicast packet and uses the forwarding class to queue mapping to decide which ingress multipoint queue forwards the packet.

When multicast path management is enabled, the ingress forwarding plane allows IP multicast snooped or routed packets to be placed on to the two multicast paths independently of the ingress classification rules. The high priority multicast path is treated as the primary path and the low priority multicast path is treated as the secondary path. The ingress bandwidth manager evaluates each multicast FIB (M-FIB) record to determine which path is best based on ingress bandwidth, number of switch fabric destinations and the fill level of each path. Explicit path association is also supported.

Dynamic Bandwidth Activity Monitoring

When ingress multicast path management is enabled on an MDA, the system monitors the in-use bandwidth associated with each Layer 2 and Layer 3 ingress multicast record. When records are first populated by static, snooping or routing protocols, they are first assumed to be inactive. An inactive record is not considered to be currently consuming ingress multicast path bandwidth.

Within the multicast-info-policy, the bandwidth activity of the new record was configured to be either managed based on an administrative bandwidth, or based on the dynamic bandwidth rate table. The bandwidth-policy associated with ingress MDA contains the configuration parameters for creating the dynamic bandwidth rate table. The purpose of the table is to allow for the system to monitor the bandwidth activity associated with a multicast record and compare the current rate against a number of rate thresholds. Rate thresholds are used to allow a multicast streams rate to fluctuate between a given range while keeping the managed rate at a certain level. Multiple dynamic managed rates are supported in the table to allow monitoring of different types of multicast traffic. Each rate threshold is associated with a rising and falling threshold that defines when the specified rate should be used and when the next lower rate should be used.

Once a record’s monitored current rate rises to the first dynamic rising threshold, the record is considered to be active and the system then manages the bandwidth the record represents based on the parameters associated with the record in the records multicast-info-policy and the configured path information in the MDAs associated bandwidth-policy.

Ingress Multicast Path Parameters

The bandwidth-policy also contains the configuration parameters for each of the managed ingress multicast paths. The queue default parameters can be overridden for each primary and secondary path. In addition, the number of secondary paths (and by implication the number of primary paths) can be overridden.

Default Bandwidth Policy

A bandwidth policy with the name 'default' always exists and is used as the default bandwidth policy when ingress multicast path management is enabled without an explicit bandwidth policy defined on an FP. The default policy cannot be deleted or edited.

The **no** form of the command removes the specified bandwidth policy from the system. The bandwidth policy associations must be removed from MDA configurations before it can be removed.

Default bandwidth-policy "default"

Parameters *policy-name* — Specifies the name of the bandwidth policy, up to 32 characters. Each bandwidth policy must be uniquely named within the system. 32 policies can be configured per system.

create — This keyword is required if creating a new bandwidth policy when the system is configured to require the explicit use of the keyword to prevent accidental object creation. Objects may be accidentally created when this protection is disabled and an object name is mistyped when attempting to edit the object. This keyword is not required when the protection is disabled. The keyword is ignored when the bandwidth policy name already exists.

admin-bw-threshold

Syntax **admin-bw-threshold** *kilo-bits-per-second*
no admin-bw-threshold

Context config>mcast-mgmt>bw-plcy

Description This command defines at which bandwidth rate a multicast channel configured to use an administrative rate starts and stop using that rate as the in-use ingress bandwidth when managing ingress multicast paths. This parameter only applies to channels that are configured to use the admin-bw rate with the **bw-activity use-admin-bw** command (both are configured in the multicast-info-policy associated with the channel context).

To be effective, the **admin-bw-threshold** value must be less than the channels configured admin-bw. If the administrative bandwidth configured on the channel is less than the administrative bandwidth threshold defined in the bandwidth policy, the admin-bw value is ignored for ingress multicast path management and the system continually uses the dynamic ingress bandwidth associated with the channel. Since the value is defined in the bandwidth-policy and the channel admin-bw value is defined in the multicast-info-policy, it is not possible to pre-determine that a given administrative bandwidth value is less than an administrative

bandwidth threshold. Since a typical administrative bandwidth threshold is set significantly lower than any administrative bandwidth values, this corner case is not expected to be prevalent. However, if the case does arise in a production environment, no ill behavior is expected as the threshold is simply a tuning parameter used to detect when the bandwidth associated with a channel has risen above any OAM or background type traffic.

While a channel that is configured to use-admin-bw (in the bw-activity command) current bandwidth is less than the admin-bw-threshold, the system treats the channel as a dynamic type channel. Once the threshold is crossed, the system immediately allocates the full admin-bw value to the channel and manages the ingress multicast path accordingly. If the bandwidth monitored on the channel rises above the admin-bw value, the system reverts to dynamic bandwidth management operation. If the bandwidth drops below the admin-bw value, but is above the admin-bw-threshold, the system uses the admin-bw value. If the bandwidth drops below the admin-bw-threshold, the system goes back to dynamic bandwidth management operation.

This command has no effect on multicast ECMP or egress CAC management operations.

The **no** form of the command reverts to the default, which is 10 kb/s.

Parameters *kilobits-per-second* — Specifies the defines the rate at which channels configured to use administrative bandwidths change from dynamic bandwidth management to using the channels configured administrative bandwidth. The parameter is expressed as an integer value and represents multiples of 1,000 bits per second. A value of 3000 indicates 3,000,000 bits per second.

Values 1 to 40,000,000

t2-paths

Syntax **t2-paths**

Context config>mcast-mgmt>bw-plcy

Description This command enables the context to configure T2 path queuing parameters for primary and secondary paths.

primary-path

Syntax **primary-path**

Context config>mcast-mgmt>bw-plcy>t2-paths

Description This command enables the context to configure primary path parameters.

secondary-path

Syntax **secondary-path**

Context config>mcast-mgmt>bw-plcy>t2-paths

Description This command enables the context to configure secondary path queue parameters. This command overrides the default path limit for the secondary path, which is one of the three ingress multicast paths into the switch fabric.

number-paths

Syntax **number-paths** *number-of-paths* [**redundant-sfm** *number-of-paths*]

Context config>mcast-mgmt>bw-plcy>t2-paths>secondary-paths

Description This command is used to explicitly provision the number of secondary paths (and imply the number of primary paths) supported by the TChip based forwarding plane the bandwidth policy is managing. The default (and minimum) number of secondary paths is 1 and the maximum configurable is 15. The number of primary paths is total number of available paths minus the number of secondary paths.

Secondary paths are used by:

- Expedited VPLS, IES and VPRN service ingress multipoint queues
- Expedited network ingress multipoint queues
- Managed multicast explicit path primary channels (using the primary paths managed multipoint queue)
- All managed multicast dynamic path channels when the primary paths or multicast planes are not at their limit (using the primary paths managed multipoint queue)
- Highest preference managed multicast dynamic path channels when the primary paths or multicast planes are at their limit (using the primary paths managed multipoint queue)

Secondary paths are used by:

- Best-Effort VPLS, IES and VPRN service ingress multipoint queues
- Best-Effort network ingress multipoint queues
- Managed multicast explicit path secondary channels (using the secondary paths managed multipoint queue)
- Lower preference managed multicast dynamic path channels when the primary paths or multicast planes are at their limit (using the secondary paths managed multipoint queue)

The number of secondary paths should be increased from the default value of 1 when a single secondary path is insufficient for the amount of explicit secondary path managed traffic or the amount of best-effort multipoint non-managed queue traffic.

The **no** form of the command restores the default number of secondary paths.

Default number-paths 1 redundant-sfm 1

Parameters *number-of-paths* — Specifies the number of secondary paths when only one switch fabric is active, while the dual-sfm parameter specifies the same value when two switch fabrics are active.

Values 1 to 15

Default 1

queue-parameters

Syntax **queue-parameters**

Context config>mcast-mgmt>bw-plcy>t2-paths>primary-path
config>mcast-mgmt>bw-plcy>t2-paths>secondary-path

Description This command enables the context to configure queue parameters. This command defines the individual parameters for the queues through which multicast packets are forwarded into the switch fabric on each path.

The individual path queues may be viewed as shared queues. All multicast packets forwarded through the switch fabric associated with one of the paths traverses bypass the normal queuing behavior. Instead of being forwarded through the normal service or network multicast queue, a single queue associated with the multicast path is used. In order to retain billing and diagnostic information, the forwarding and discard statistics for the service or network queue the packet would have traversed without ingress multicast management is used to account for each packets behavior.

Note: Any ingress scheduling policy functions attempting to manage the service or network multicast queues is only able to read the statistics of the multicast queues and not able to manage the queues dynamic rate since the packets are flowing through different, non-managed queues. Since this is the case, multicast queues parented to a scheduling policy should be parented to the root scheduler at the highest priority without any rate limitation. Any ingress rate limiting for multicast traffic is performed by the multicast path bandwidth manger based on each records priority and a possible “black-hole” rate threshold.

All queues created for ingress multicast path management are automatically created by the system out of the system reserved queue space. Each queue is created as an expedited queue.

When forwarding through the queues, each packets forwarding class is ignored. However, the forwarding class is retained for proper egress processing. The packets expressed or implied profile is also ignored within the ingress path queues. A packets congestion priority is derived from the records cong-priority-threshold evaluation result as indicated by the policy or multicast-info-policy. The cong-priority-threshold sets the high or low congestion priority of a record based on the records preference value. Within each multicast information policy bundle the cong-priority-threshold is set with a value from 0 to 7 and defines the threshold at

which all records with a preference equal to or higher than the defined preference is treated as congestion priority high. Multicast records with a preference lower than the defined class threshold is treated as congestion priority low. Low-priority packets use the low drop-tail threshold of the queue, while high-priority packets use the standard MBS value. In the event of path congestion, low-priority packets are discarded first, leaving room for the higher priority packets.

For the primary and secondary paths, a single queue exists for each path and every packet forwarded through the path by the bandwidth manager uses that queue.

cbs

Syntax **cbs** *percent-of-resv-cbs*
no cbs

Context config>mcast-mgmt>bw-plcy>t2-paths>primary-paths>queue-parameters
config>mcast-mgmt>bw-plcy>t2-paths>secondary-paths>queue-parameters

Description This command overrides the default Committed Buffer Size (CBS) for each individual path's queue. The queues CBS threshold is used when requesting buffers from the systems ingress buffer pool to indicate whether the requested buffer should be removed from the reserved portion of the buffer pool or the shared portion. When the queue's fill depth is below or equal to the CBS threshold, the requested buffer comes from the reserved portion. Once the queues depth exceeds the CBS threshold, buffers come from the shared portion.

The **cbs** *percent-of-resv-cbs* parameter is defined as a percentage of the reserved portion of the pool. The system allows the sum of all CBS values to equal more than 100% allowing for oversubscription of the reserved portion of the pool. If the reserved portion is oversubscribed and the queues are currently using more reserved space than provisioned in the pool, the pool automatically starts using the shared portion of the pool for within-CBS buffer allocation. On the shared early detection slopes could assume more buffers exist within the shared portion than actually do which may cause the early detection function to fail.

For the primary-path and secondary-path queues, the percentage is applied to a single queue for each path.

The **no** form of the command restores the path queues default CBS value.

Parameters *percent-of-resv-cbs* — The percent of buffers reserved from the total buffer pool space, expressed as a decimal integer. If 10 MB is the total buffers in the buffer pool, a value of 10 would reserve 1MB (10%) of buffer space for the forwarding class queue. The value 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can be applied for scheduling purposes).

Values 0 to 100

Default

Primary: 5
Secondary: 30

drop-tail

Syntax **drop-tail**

Context config>mcast-mgmt>bw-plcy>t2>prim-path>queue
config>mcast-mgmt>bw-plcy>t2>sec-path>queue

Description This command enters the context to configure queue drop-tail parameters.

low

Syntax **low**

Context config>mcast-mgmt>bw-plcy>t2>prim-path>queue>drop-tail
config>mcast-mgmt>bw-plcy>t2>sec-path>queue>drop-tail

Description This command enters the context to configure the queue low drop-tail parameters. The low drop-tail defines the queue depth beyond which out-of-profile packets are not accepted into the queue and are discarded.

percent-reduction-from-mbs

Syntax **percent-reduction-from-mbs** *percent*
no percent-reduction-from-mbs

Context config>mcast-mgmt>bw-plcy>t2>prim-path>queue>drop-tail
config>mcast-mgmt>bw-plcy>t2>sec-path>queue>drop-tail

Description This command overrides the default percentage value used to determine the low drop-tail value for the queue. By default, 10 percent of the queue depth is reserved for high congestion priority traffic. When specified, the **percent-reduction-from-mbs** percentage value is applied to the queues' MBS threshold. The resulting value is subtracted from the MBS to derive the low drop-tail threshold maintained by the queue. The low drop-tail threshold defines the point at which all low-congestion priority packets destined for the queue are discarded based on queue depth. Low- and high-congestion priority is derived from the multicast records preference value compared to the record's bundle priority threshold.

The **no** form of this command restores the default value.

Default 10%

Parameters *percent* — Specifies the percent of queue depth reserved for high-congestion priority traffic.

Values 0 to 100, default

0 specifies that the MBS and low drop-tail thresholds be set to the same value resulting in high- and low-congestion priority packets being treated equally.

100 specifies that the low drop-tail threshold is set to 0, resulting in all low-congestion priority packets being discarded.

Values between 0 and 100 result in a corresponding differential between the MBS and the low drop-tail threshold values.

mbs

Syntax **mbs** *percent-of-pool*
no mbs

Context config>mcast-mgmt>bw-plcy>t2-paths>primary-paths>queue-parameters
config>mcast-mgmt>bw-plcy>t2-paths>secondary-paths>queue-parameters

Description This command configures the override for the default Maximum Buffer Size (MBS) for each individual path's queue. The queues MBS threshold defines the point at which all packets destined for the queue are discarded based on queue depth. The defined threshold also provides context for the queues drop-tail parameter.

The mbs *percent-of-pool* parameter is defined as a percentage of the total pool size. The system allows the sum of all MBS values to equal more than 100% allowing for oversubscription of the pool.

For the primary-path and secondary-path queues, the mbs percent is applied to a single queue for each path.

The **no** form of the command is used to restore the path queues default MBS value.

Parameters *percent-of-pool* — The percent of buffers from the total buffer pool space for the maximum amount of buffers, expressed as a decimal integer. If 10 MB is the total buffers in the buffer pool, a value of 10 would limit the maximum queue size to 1MB (10%) of buffer space for the forwarding class queue. If the total size is increased to 20MB, the existing value of 10 would automatically increase the maximum size of the queue to 2MB.

Values 0 to 100

Default

Primary:	7
Secondary:	40

falling-percent-reset

Syntax **falling-percent-reset** *percent-of-highest*
no falling-percent-reset

Context config>mcast-mgmt>bw-plcy

Description This command is configures the percentage of bandwidth decrease that must occur to reset the dynamic bandwidth monitoring function for a multicast channel. When a channel is configured to use the ingress dynamic bandwidth as the in-use bandwidth for ingress multicast path management, the system maintains a sliding window in time that defines how long the last highest bandwidth value associated with the channel should be used. The sliding window duration is derived from the channels bw-activity dynamic falling-delay parameter within the multicast information policy. Each time the system detects a current bandwidth for a channel that is equal to or greater than the current highest bandwidth for the channel, the sliding window is reset and the highest value is used when managing the ingress multicast paths. If the system does not detect a higher or equal bandwidth value for the channel within the window period, the system resets the sliding window and uses the next highest rate seen during the duration of the window period. In this way, the system delays relinquishing bandwidth for a dynamic bandwidth channel for a configurable period of time. If a momentary fluctuation (decrease) in ingress bandwidth occurs, the system ignores the bandwidth change.

While this is useful for momentary fluctuations in bandwidth, it may be desirable to react faster when the current bandwidth monitored for a channel drops significantly relative to the currently in-use bandwidth. When the bandwidth decrease is equal to or greater than the falling-percent-reset value, the system immediately stops using the highest bandwidth and starts using the current bandwidth while resetting the sliding window.

If falling-percent-reset is set to 50%, when the current ingress dynamic bandwidth is 50% of the current in-use highest bandwidth, the system immediately uses the current dynamic ingress bandwidth as the highest bandwidth for the channel.

By default falling-percent-reset is 50% when a new bandwidth policy is created. The default bandwidth policy also has a hard configured value of 50%. Setting falling-percent-reset to 100 is equivalent to specifying no falling-percent-reset.

The **no** form of the command restores the default value of 50%.

Parameters *percent-of-highest* — The percent-of-highest parameter is required and defines the percentage of decline between the current ingress dynamic bandwidth and the current in-use highest bandwidth at which the system resets the dynamic ingress bandwidth monitoring for the channel. When reset in this case, the system uses the current ingress dynamic bandwidth as the highest rate and continues monitoring. The parameter must be defined as an integer value representing a percentage.

Values 1 to 100 percent

Default 50

mcast-pool

Syntax **mcast-pool percent-of-total** *percent-of-buffers* **resv-cbs** *percent-of-pool* **slope-policy** *policy-name*

no mcast-pool

Context config>mcast-mgmt>bw-plcy

Description This command configures the ingress multicast path management buffer pool. The pool is used by the primary and secondary path queues through which all ingress managed multicast traffic must flow. The parameters may be used to configure the size of the pool relative to the total ingress buffer space, the amount of reserved CBS buffers within the pool and the slope policy used to manage early congestion detection functions in the shared portion of the pool.

Care should be taken when managing the buffer pool space as changes to the systems buffer pool behavior can have negative effects on multicast and unicast forwarding.

Sizing the Pool

The percent-of-total command defines how much of the total ingress buffer pool space for the MDA is dedicated for multicast channels managed by the bandwidth policy. Since multicast typically has a higher scheduling priority through the switch fabric, the buffer pool does not need to be large. By default, the system reserves 10% of the buffers on the ingress side of the MDA once multicast path management is enabled.

Reserved CBS Portion of the Pool

The multicast pool is divided into two portions; reserved and shared. The reserved portion is used by the multicast path queues until they cross there individual CBS thresholds. Since the CBS thresholds are configured as percents and the percents are allowed to oversubscribe the reserved portion of the pool, it is possible for some of the queues CBS buffer allocation to be met by the shared portion of the pool. By default, 50% of the pool is defined as reserved. This may be changed using the resv-cbs percentage parameter.

Shared Portion WRED Slopes

The shared portion of the buffer pool is used by queues that have crossed over their CBS thresholds. Since the total MBS values for the multicast path queues may oversubscribe the pool size, a buffer congestion control mechanism is supported within the pool in the form of two WRED slopes. The slope-policy parameter defines how the slopes are configured and whether they are activated. Each packet entering a path queue is defined as high or low priority within the queue based on the channels preference value relative to the cong-priority-threshold command. When getting a shared buffer of a high priority packet, the high WRED slope is used. Low priority packets use the low WRED slope.

The **no** form of the command returns the managed multicast path pool to its default settings.

Parameters *percent-of-buffers* — The percent-of-total keyword is required when executing the mcast-pool command and must be followed by a percent-of-buffers parameter expressed as an integer and representing the percentage of ingress buffers that is allocated to the multicast pool.

Values 1 to 50

Default 10

percent-of-pool — The resv-cbs keyword is required when executing the mcast-pool command and must be followed by a percent-of-pool parameter expressed as an integer and representing the percentage of the pool that is reserved for multicast path queues within their CBS threshold.

Values 1 to 100

Default 50

slope-policy-name — The slope-policy keyword is required when executing the mcast-pool command and must be followed by a valid slope-policy-name. The named policy is used to configure the WRED slopes within the multicast pool. Once a slope policy is associated with a buffer pool, it cannot be deleted.

Default default

8.6.2.1.4 Multicast Info Policy Commands

multicast-info-policy

Syntax **multicast-info-policy** *policy-name* [**create**]
no multicast-info-policy

Context config>mcast-management

Description This command configures a multicast information policy. Multicast information policies are used to manage parameters associated with Layer 2 and Layer 3 multicast records. Multiple features use the configured information within the policy. The multicast ingress path manager uses the policy to decide the inactive and active state behavior for each multicast record using the ingress paths to the switch fabric. The system's multicast ECMP join decisions are influenced by the channel information contained within the policy.

Multicast Bundles:

A multicast information policy consists of one or multiple named bundles. Multicast streams are mapped to a bundle based on matching the destination address of the multicast stream to configured channel ranges defined within the bundles. Each policy has a bundle named 'default' that is used when a destination address does not fall within any of the configured channel ranges.

Each bundle has a set of default parameters used as the starting point for multicast channels matching the bundle. The default parameters may be overridden by optional exception parameters defined under each channel range. Further optional parameter overrides are possible under explicit source address contexts within each channel range.

Default Multicast Information Policy

A multicast information policy always exists with the name 'default' and cannot be edited or deleted. The following parameters are contained in the default multicast information policy:

Policy Description:	Default policy, cannot be edited or deleted.
Bundle:	default
Bundle Description:	Default Bundle, cannot be edited or deleted.
Congestion-Priority-Threshold:	4
ECMP-Optimization-Limit-Threshold:	7
Bundle Defaults:	
Administrative Bandwidth:	0 (undefined)
Preference:	0
CAC-Type:	Optional
Bandwidth Activity:	Dynamic with no black-hole rate
Explicit Ingress SF Path:	None (undefined)
Configured Channel Ranges:	None

The default multicast information policy is applied to all VPLS and VPRN services and all routing contexts until an explicitly defined multicast information policy has been mapped.

Explicit Multicast Information Policy Associations

Each VPLS service and each routing context (including VPRN routing contexts) supports an explicit association with an pre-existing multicast information policy. The policy may need to be unique per service or routing context due to the fact that each context has its own multicast address space. The same multicast channels may be and most likely be used for completely different multicast streams and applications in each forwarding context.

Interaction with Ingress Multicast Path Management

When ingress multicast path management is enabled on an MDA, the system automatically creates a bandwidth manager context that manages the multicast path bandwidth into the switch fabric used by the ingress ports on the MDA. As routing or snooping protocols generate Layer 2 or Layer 3 multicast FIB records that are populated on the MDA's forwarding plane, they are processed though the multicast information policy that is associated with the service or routing context associated with the record. The policy returns the following information for the record to be used by the ingress bandwidth manager:

- The records administrative bandwidth (0 if undefined)
- Preference level (0 to 7 with 7 being highest)
- Bandwidth activity monitoring setting (use admin bw or dynamic monitoring)
  If admin bw is indicated, also returns active and inactive thresholds
- Initial switch fabric multicast path (primary or secondary)
- Explicit switch fabric multicast path (primary, secondary, or none)

Interaction with Multicast ECMP Optimization

The multicast information policy is used by the multicast ECMP optimization function to derive each channels administrative bandwidth. The ECMP function tallies all bandwidth information for channels joined and attempts to equalize the load between the various paths to the sender. The multicast information policy returns the following information to the ECMP path manager:

- Administrative bandwidth (0 if undefined)
- Preference (0 to 7 with 7 the highest preference value)

Default multicast-info-policy “default”

Parameters *policy-name* — Identifies the name of the policy to be either created or edited. Each multicast information policy must be uniquely named within the system. Names of up to 32 ASCII characters are supported with the normal character restrictions.

create — The create keyword is required if creating a new multicast information policy when the system is configured to require the explicit use of the keyword to prevent accidental object creation. Objects may be accidentally created when this protection is disabled and an object name is mistyped when attempting to edit the object. This keyword is not required when the protection is disabled. The keyword is ignored when the multicast information policy name already exists.

multicast-info-policy

Syntax **multicast-info-policy** *policy-name*
no multicast-info-policy

Context config>service>vpls
config>router

Description This command overrides the default multicast information policy on a VPLS or routing context. When the policy association is changed, all multicast channels in the service or routing context must be reevaluated.

If a multicast information policy is not explicitly associated with the VPLS service or routing context, the default multicast information policy is used when ingress multicast path management is enabled.

While a multicast information policy is associated with a service or routing context, the policy cannot be deleted from the system.

The **no** form of the command removes an explicit multicast information policy from the VPLS or routing context and restores the default multicast information policy.

Parameters *policy-name* — This parameter is required and specifies an existing multicast information policy that should be associated with the VPLS service or routing context.

Default default

bundle

Syntax **bundle** *bundle-name* [**create**]
no bundle *bundle-name*

Context config>mcast-mgmt>mcast-info-plcy

Description This command creates or edits channel bundles within a multicast information policy. Bundles are used for two main purposes. First, bundles are used by the multicast CAC function to group multicast channels into a common bandwidth context. The CAC function limits the ability for downstream nodes to join multicast channels based on the egress interfaces ability to handle the multicast traffic. Bundling allows multicast channels with common preference or application to be managed into a certain percentage of the available bandwidth.

The second function of bundles is to provide a simple provisioning mechanism. Each bundle within a multicast information policy has a set of default channel parameters. If each channel provisioned in to the bundle is able to use the default parameters for the bundle, the provisioning and configuration storage requirements are minimized.

Up to 31 explicit bundles may be defined within a multicast information policy (32 including the default bundle).

Once a bundle is created, the default channel parameters should be configured and the individual channel ranges should be defined. Within each channel range, override parameters may be defined that override the default channel parameters. Further overrides are supported within the channel range based on explicit source overrides.

A bundle may be deleted at anytime (except for the default bundle). When a bundle is deleted, all configuration information within the bundle is removed including multicast channel ranges. Any multicast records using the bundle should be reevaluated. Multicast CAC and ECMP managers should also be updated.

Default Bundle

Each multicast information policy contains a bundle named **default**. The default bundle cannot be deleted. Any multicast channel that fails to match a channel range within an explicit bundle is automatically associated with the default bundle.

The **no** form of the command removes a bundle from the multicast information policy. The default bundle cannot be removed from the policy.

Default bundle “default”

Parameters *bundle-name* — Specifies bundle expressed as an ASCII string with up to 16 characters and must follow normal naming conventions. If bundle-name already exists, the system enters the bundle context for editing purposes. If bundle-name does not exist, the system creates the defined bundle in the policy and enter the bundle context for editing purposes.

create — The create keyword is required if creating a new multicast information policy bundle when the system is configured to require the explicit use of the keyword to prevent accidental object creation. Objects may be accidentally created when this protection is disabled and an object name is mistyped when attempting to edit the object. This keyword is not required when the protection is disabled. The keyword is ignored when the bundle name already exists.

channel

Syntax **channel** *ip-address* [*ip-address*] [**create**]
no channel *ip-address* [*ip-address*]

Context config>mcast-mgmt>mcast-info-plcy>bundle

Description This command defines explicit channels or channel ranges that are associated with the containing bundle. A channel or channel range is defined by their destination IP addresses. A channel may be defined using either IPv4 or IPv6 addresses. If a channel range is being defined, both the start and ending addresses must be the same type.

A specific channel may only be defined within a single channel or channel range within the multicast information policy. A defined channel range cannot overlap with an existing channel range.

If a channel range is to be shortened, extended, split or moved to another bundle, it must first be removed from its existing bundle.

Each specified channel range creates a containing context for any override parameters for the channel range. By default, no override parameters exist.

The **no** form of the command removes the specified multicast channel from the containing bundle.

Parameters *ip-address* [*ip-address*] — These parameters define the starting and ending destination IP addresses for a channel range.

If only the start channel *ip-address* parameter is given, the channel ranges comprises of a single multicast channel.

If both the starting and ending address are specified, all addresses within the range including the specified address are part of the channel range.

IPv4 or IPv6 addresses may be defined. All specified addresses must be valid multicast destination addresses. The starting IP address must be numerically lower then the ending IP address.

Values Any valid IP multicast destination address

create — The create keyword is required if creating a new multicast channel range when the system is configured to require the explicit use of the keyword to prevent accidental object creation. Objects may be accidentally created when this protection is disabled and an object name is mistyped when attempting to edit the object. This keyword is not required when the protection is disabled. The keyword is ignored when the specified channel range already exists.

admin-bw

Syntax **admin-bw** *kbps*
no admin-bw

Context config>mcast-mgmt>mcast-info-plcy>bundle>channel
config>mcast-mgmt>mcast-info-plcy>bundle>source-override

Description This command specifies an administrative bandwidth for multicast channels.

The specified bandwidth rate may be used by the multicast ingress path manger, multicast CAC manager or multicast ECMP manager.

The *kbps* value is closely tied to the **bw-activity** command. When the bw-activity command is set to use-admin-bw, the multicast ingress path manager uses the configured administrative bandwidth value as the managed ingress bandwidth. The **admin-bw** value must be defined for the **bw-activity use-admin-bw** command to succeed. Once the bw-activity command is set to use the admin-bw value, the value cannot be set to 0 and the no admin-bw command fails. Setting the bw-activity command to dynamic (the default setting), breaks the association between the commands.

The **no** form of the command restores the default value for admin-bw. If the command is executed in the **channel** context, the channels administrative bandwidth value is set to null. If the command is executed in the **source-override** context, the source override administrative bandwidth value is set to null.

Parameters *kbps* — Specifies the administrative bandwidth for multicast channels.

Values 1 to 40000000 kb/s

Bundle default:	0
Channel default:	Null (undefined)
Source-override default:	Null (undefined)

Override sequence — The channel setting overrides the bundle setting. The source-override setting overrides the channel and bundle settings.

bw-activity

Syntax **bw-activity** {**use-admin-bw** | **dynamic** [**falling-delay** *seconds*]} [**black-hole-rate** *kbps*]

no bw-activity

Context config>mcast-mgmt>mcast-info-plcy>bundle>channel
config>mcast-mgmt>mcast-info-plcy>bundle>source-override

Description This command defines how the multicast ingress path manager determines the amount of bandwidth required by a multicast channel. The default setting is dynamic which causes the bandwidth manager to use the bandwidth policies dynamic rate table entries to determine the current rate. The alternative setting is use-admin-bw which causes the bandwidth manager to use the configured admin-bw associated with the channel. The use-admin-bw setting also requires an active and inactive threshold to be defined which allows the bandwidth manager to determine when the channel is actively using ingress path bandwidth and when the channel is idle.

The **use-admin-bw** setting requires that the channel be configured with an admin-bw value that is not equal to 0 in the same context as the **bw-activity** command using the setting. Once a context has use-admin-bw configured, the context's admin-bw value cannot be set to 0 and the no admin-bw command fails.

This command also supports an optional **black-hole-rate** *kbps* command that defines at which current rate a channel should be placed in the black-hole state. This is intended to provide a protection mechanism against multicast channels that exceed a reasonable rate and cause outages in other channels.

The **no** form of the command restores the default bandwidth activity monitoring setting (dynamic or null depending on the context).

Parameters **use-admin-bw | dynamic** — The **use-admin-bw** and **dynamic** keywords are mutually exclusive and one must be specified when executing the **bw-activity** command. The **use-admin-bw** keyword indicates the channels current ingress bandwidth should be derived from the **admin-bw** setting. The **admin-bw** setting must not currently be set to 0 for the **use-admin-bw** setting to succeed.
The **dynamic** keyword indicates that the multicast ingress path manager should use the dynamic rate table (as defined in the bandwidth-policy) to derive the channels current ingress rate.

falling-delay *seconds* — Specifies the value the bandwidth manager uses the falling-delay threshold to hold on to the previous highest bandwidth until the delay time has expired while operating in dynamic bandwidth mode. This allows the bandwidth manager to ignore momentary drops in channel bandwidth.

Values 10 to 3600

Default 30

Bundle default:	dynamic
Channel default:	Null (undefined)
Source-override default:	Null (undefined)

black-hole-rate *kbps* — Specifies a rate at which a channel is placed in the black-hole state. The kilobits-per-second parameter is expressed as an integer and represents multiples of 1,000 bits per second.

Values 1 to 40000000

Default None

explicit-sf-path

Syntax **explicit-sf-path** {**primary** | **secondary**}
no explicit-sf-path

Context config>mcast-mgmt>mcast-info-plcy>bundle>channel
config>mcast-mgmt>mcast-info-plcy>bundle>source-override

Description This command defines an explicit ingress switch fabric multicast path assigned to a multicast channel. When defined, the channel is setup with the explicit path as its inactive path. When an explicit path is not defined, all multicast channels are initialized on the secondary path and when they start to consume bandwidth, they are moved to the appropriate path based on the channel attributes and path limitations. Explicit path channels are not allowed to move from their defined path.

The **explicit-sf-path** command in the bundle context defines the initial path for all channels associated with the bundle unless the channel has an overriding explicit-sw-path defined in the channel context. The channel context may also be overridden by the explicit-sf-path command in the source-override context. The channel and source-override explicit-sf-path settings default to null (undefined) and have no effect unless explicitly set.

The **no** form of the command restores default path association behavior (dynamic or null depending on the context).

Parameters **primary** — The **primary** and **secondary** keywords are mutually exclusive to one another. One keyword must be specified when executing the explicit-sf-path command. The **primary** keyword specifies that the primary ingress multicast path should be used as the explicit path for the channel.

secondary — The **primary** and **secondary** keywords are mutually exclusive to one another. One keyword must be specified when executing the explicit-sf-path command. The **secondary** keyword specifies that the secondary ingress multicast path should be used as the explicit path for the channel.

Override sequence — The channel setting overrides the bundle setting. The source-override setting overrides the channel and bundle settings.

keepalive-override

Syntax **keepalive-override** *keepalive-timer*
no keepalive-override

Context config>mcast-mgmt>mcast-info-policy>bundle
config>mcast-mgmt>mcast-info-policy>bundle>channel
config>mcast-mgmt>mcast-info-policy>bundle>channel>source-override

Description This command configures the keepalive timer override. The PIM (S,G) Keepalive Timer (KAT) is used to maintain the (S,G) state when (S,G) join is not received. Expiry of the KAT causes the (S,G) entry to be removed.

The KAT override configuration is performed with an multicast information policy, which must be applied to the related PIM routing instance. When a KAT override is configured under a channel (a group or a group range), it applies to all (S,G) entries that fall under it, except when the source-override is configured and a KAT override is also configured under the source-override. In this scenario, the specific KAT override must be used for the (S,G) entries that fall under the source-override, while other (S,G) entries under the bundle use the KAT override configured under the channel.

Parameters *keepalive-timer* — Specifies the keepalive timer override, in seconds.

Values 10 to 300

preference

Syntax **preference** *preference-level*
no preference

Context config>mcast-mgmt>mcast-info-plcy>bundle
config>mcast-mgmt>mcast-info-plcy>bundle>channel
config>mcast-mgmt>mcast-info-plcy>bundle>channel>source-override

Description This command sets the relative preference level for multicast channels. The preference of a channel specifies its relative importance over other multicast channels. Eight levels of preference are supported; 0 through 7. Preference value 7 indicates the highest preference level.

When the multicast ingress path manager is congested on one or more of the switch fabric multicast paths, it uses the preference values associated with each multicast record to determine which records are allowed on the path and which records be placed in a black-hole state.

The preference value is also compared to the bundles cong-priority-threshold setting to determine the congestion priority of the channel. The result also dictates the channels multicast CAC class level (high or low). When the channels preference value is less than the congestion priority threshold, it is considered to have a congestion priority and CAC class value equal to low. When the channels preference value is equal to or greater than the threshold, it is considered to have a congestion priority and a CAC class value equal to high.

The preference value is also compared to the bundles ecmp-opt-threshold setting to determine whether the channel is eligible for ECMP path dynamic optimization. If the preference value is equal to or less than the threshold, the channel may be optimized. If the preference value is greater than the threshold, the channel will not be dynamically optimized.

The preference command may be executed in three contexts; bundle, channel and source-override. The bundle default preference value is 0. The channel and source-override preference settings are considered overrides to the bundle setting and have a default value of null (undefined).

The **no** form of the command restores the default preference value (0 or null depending on the context).

Parameters *preference-level* — The preference-level parameter is required and defines the preference value of the channel.

Values 1 to 7

Bundle default:	0
Channel default:	Null (undefined)
Source-override default:	Null (undefined)

Override sequence — The channel setting overrides the bundle setting. The source-override setting overrides the channel and bundle settings.

primary-tunnel-interface

Syntax **primary-tunnel-interface** {**rsvp-p2mp** *lsp-name* | **ldp-p2mp** *p2mp-id*} **sender** *ip-address*
no primary-tunnel-interface

Context config>mcast-mgmt>mcast-info-plcy>bundle
config>mcast-mgmt>mcast-info-plcy>bundle>channel
config>mcast-mgmt>mcast-info-plcy>bundle>channel>source-override

Description This command allows the user to define a bundle in the multicast-info-policy and specify channels in the bundle that must be received from the primary tunnel interface associated with an RSVP P2MP LSP. The multicast info policy is applied to the base router instance.

The egress LER can accept multicast packets via two different methods. The regular RPF check on unlabeled IP multicast packets, which is based on routing table lookup. The static assignment which specifies the receiving of a multicast group <*,G> or a specific <S,G> from a primary tunnel-interface associated with an RSVP P2MP LSP.

One or more primary tunnel interfaces in the base router instance can be configured. That is, the user can specify to receive different multicast groups, <*,G> or specific <S,G>, from different P2MP LSPs. This assumes that there are static joins configured for the same multicast groups at the ingress LER to forward over a tunnel interface associated with the same P2MP LSP.

At any given time, packets of the same multicast group can be accepted from either the primary tunnel interface associated with a P2MP LSP or from a PIM interface. These are mutually exclusive options. As soon as a multicast group is configured against a primary tunnel interface in the multicast info policy, it is blocked from other PIM interfaces.

A multicast packet received on a tunnel interface associated with a P2MP LSP can be forwarded over a PIM or IGMP interface which can be an IES interface, a spoke SDP terminated IES interface, or a network interface.

The **no** form of the command removes the static RPF check.

Parameters *lsp-name* — Species a string of up to 32 characters identifying the LSP name as configured at the ingress LER.

ip-address — Specifies a string of 15 characters representing the IP address of the ingress LER for the LSP.

p2mp-id — Identifier used for signaling mLDP P2MP LSP (applies only to the 7750 SR).

Values 1 to 4294967296

source-override

Syntax **source-override** *ip-address* [**create**]
no source-override *ip-address*

Context config>mcast-mgmt>mcast-info-plcy>bundle>channel

Description This command defines a multicast channel parameter override context for a specific multicast sender within the channel range. The specified senders IP address must be of the same type (IPv4 or IPv6) as the containing channel range.

The **no** form of the command removes the specified sender override context from the channel range.

Parameters *ip-address* — Specifies either am IPv4 or IPv6 address and it must be the same type as the containing channel range.

Values

ipv4-address	a.b.c.d
ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D

create — The **create** keyword is required if creating a new source override when the system is configured to require the explicit use of the keyword to prevent accidental object creation. Objects may be accidentally created when this protection is disabled and an object name is mistyped when attempting to edit the object. This keyword is not required when the protection is disabled. The keyword is ignored when the specified source override IP address already exists.

cong-priority-threshold

Syntax **cong-priority-threshold** *preference-level*
no cong-priority-threshold

Context config>mcast-mgmt>mcast-info-plcy>bundle

Description This command defines the preference level threshold where records change from low congestion priority to high congestion priority. Congestion priority is used by the ingress multicast path queues to map packets entering the queue to either the low drop-tail or the MBS drop-tail threshold. In the event that congestion happens on the queue, the queue depth increases. As the queue depth increases beyond the low drop-tail, packets with low priority congestion priority are discarded. This leaves room in the queue for packets with high congestion priority until the queue reaches the MBS threshold.

The default congestion priority threshold is 4. This means that multicast channels with a preference level of 0 to 3 are treated as having low congestion priority and channels with preference level of 4 to 7 are treated as having a high congestion priority. The **cong-priority-threshold** command can be used to change the default threshold. Any multicast channel with a preference equal to or higher than the configured threshold is treated with high congestion priority.

The **cong-priority-threshold** value is also used by the multicast CAC manager to derive the class of a channel matched by the multicast information policy. Channels with a preference less than the configured threshold are treated as low class and channels with a preference equal to or greater than the threshold is treated as high class.

Changing the **cong-priority-threshold** value causes all channels congestion priority to be reevaluated. Both the ingress multicast path managers and multicast CAC managers must be updated.

The **no** form of the command restores the default congestion priority preference threshold value.

Parameters *preference-level* — The preference-level parameter is required when specifying the cong-priority-threshold.

Values 0 to 7

Default 4

ecmp-opt-threshold

Syntax **ecmp-opt-threshold** *preference-level*
no ecmp-opt-threshold

Context config>mcast-mgmt>mcast-info-plcy>bundle

Description This command defines the preference level threshold where multicast ECMP path management is allowed to dynamically optimize channels based on topology or bandwidth events. If the channels preference is equal to or less than the ecmp-opt-threshold, ECMP is allowed to move the channel between ECMP paths when bandwidth distribution events happen. Channels with a preference level higher than the threshold are moved during these events.

The default ECMP optimization limit threshold is 7. This means that multicast channels with a preference level of 0 to 7 (all channels) are allowed to move between ECMP paths. The ecmp-opt-threshold command can be used to change the default threshold.

Changing the threshold causes all channels ECMP optimization eligibility to be reevaluated.

The **no** form of the command restores the default ECMP optimization preference threshold value.

Parameters *preference-level* — The preference-level parameter is required when specifying the ecmp-opt-threshold. An integer value from 0 to 7 must be specified.

Values 0 to 7

Default 7

mc-ecmp-balance

Syntax [**no**] **mc-ecmp-balance**

Context config>router>pim

Description This command enables multicast balancing of traffic over ECMP links considering multicast bandwidth. When enabled, every multicast stream that needs to be forwarded over an ECMP link is evaluated for the sum total multicast utilization currently using the ECMP interface in question.

Note: A given interface can be shared between multiple (partially overlapping) ECMP sets. This is taken into consideration and a complete balance is attempted.

ECMP load balancing helps to avoid loss of unicast traffic over ECMP links as it loads balance over ECMP links and if multicast is not balanced then it is possible that a given link does not have sufficient bandwidth to pass its allotted unicast traffic portion.

In order to achieve a proper balance, multicast groups and their bandwidth should be configured in the config **mcast-management** context.

If the bandwidth is not configured, then the default value applies, and for the purpose of ECMP load balancing, the net effect is that the balance achieved reflects a balance of the number of multicast groups traversing over the ECMP link. The bandwidth used in this policy is the configured value, not the actual bandwidth.

If **mc-ecmp-balance** is enabled, a redistribution may be triggered whenever an interface is added to an ECMP link.

If **mc-ecmp-balance** is enabled, a period re-balance may be configured that re-optimizes the distribution as some multicast streams may have been removed from the ECMP link.

If **mc-ecmp-balance** is enabled, then a threshold (**ecmp-opt-threshold**) can be configured to avoid moving multicast streams where interruption should be avoided.

The ecmp-opt-threshold is used to define the preference level threshold where multicast ECMP path management is allowed to dynamically optimize channels based on topology or bandwidth events. If the channels preference is equal to or less than the ecmp-opt-threshold, ECMP is allowed to move the channel between ECMP paths when bandwidth distribution events happen. Channels with a preference level higher than the threshold are not moved during these events.

Changing the **ecmp-opt-threshold** causes all channels ECMP optimization eligibility to be reevaluated.

The **no** form of the command removes the re-balancing capability from the configuration.

Default mc-ecmp-balance

mc-ecmp-balance-hold

Syntax **mc-ecmp-balance-hold** *minutes*
no mc-ecmp-balance-hold

Context config>router>pim

Description This command defines a hold period that applies after an interface has been added to the ECMP link. It is also used periodically to rebalance if channels have been removed from the ECMP link.

If the ECMP interface has not changed in the hold period and if no multicast streams have been removed, then no action is taken when the timer triggers.

This parameter should be used to avoid excessive changes to the multicast distribution.

A rebalance will not occur to multicast streams that have a priority greater then the configured ecmp-opt-threshold.

The **no** form of the command reverts to default.

Parameters *minutes* — Specifies the hold down time in minutes.

Values 2 to 600

mc-ecmp-rebalance

Syntax **mc-ecmp-rebalance** [**ecmp-opt-threshold**]

Context tools>perform>router>pim

Description This command triggers an immediate rebalance, regardless if the hold timer has triggered or if any changes have occurred.

Parameters **ecmp-opt-threshold** — Forces a rebalance of all multicast streams with a priority equal or less then the specified value.

Specifying the value of 7 forces all multicast streams to be re-balanced regardless of the configured **ecmp-opt-threshold** value.

Values 1 to 7

8.6.2.1.5 Service Commands

mcac

Syntax **mcac**

Context config>service>pw-template>igmp-snooping
config>service>vpls>mesh-sdp>snooping

Description This command enables the context to configure multicast CAC policy parameters and constraints for this interface.

mld-snooping

Syntax **mld-snooping**

Context config>service>vpls
config>service>vpls>sap

Description This command enables the context to configure MLD snooping parameters.

policy

Syntax **policy** *policy-name*
no policy

Context config>service>pw-template>igmp-snooping>mcac
config>service>vpls>mesh-sdp>igmp-snp>mcac

Description This command configures the multicast CAC policy name.

The **no** form of the command reverts to the default value.

Parameters *policy-name* — Specifies multicast CAC policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

unconstrained-bw

Syntax **unconstrained-bw** *bandwidth* **mandatory-bw** *mandatory-bw*
no unconstrained-bw

Context config>service>vpls>mesh-sdp>igmp-snp>mcac
config>service>vpls>spoke-sdp>igmp-snp>mcac
config>service>vpls>sap>igmp-igmp-snp>mcac

Description This command configures the bandwidth for the interface's multicast CAC policy traffic. When disabled (**no unconstrained-bw**) there is no checking of bandwidth constraints on the interface level. When enabled and a policy is defined, enforcement is performed. The allocated bandwidth for optional channels should not exceed the **unconstrained-bw** minus the **mandatory-bw** and the mandatory channels have to stay below the specified value for the **mandatory-bw**. After this interface check, the bundle checks are performed.

The **no** form of the command reverts to the default value.

Parameters *bandwidth* — The bandwidth assigned for interface's MCAC policy traffic, in kilobits per second (kb/s).

Values 0 to 2147483647

mandatory-bw — Specifies the bandwidth pre-reserved for all the mandatory channels on a given interface in kilobits per second (kb/s).

If the *bandwidth* value is 0, no mandatory channels are allowed. If *bandwidth* is not configured, then all mandatory and optional channels are allowed.

If the value of *mandatory-bw* is equal to the value of *bandwidth*, then all the unconstrained bandwidth on a given interface is allocated to mandatory channels configured through multicast CAC policy on that interface and no optional groups (channels) are allowed.

The value of *mandatory-bw* should always be less than or equal to that of *bandwidth*, An attempt to set the value of *mandatory-bw* greater than that of *bandwidth*, results in inconsistent value error.

Values 0 to 2147483647

8.6.2.1.6 VPLS Multicast Commands

vpls

Syntax **vpls** *service-id* [**customer** *customer-id*] [**vpn** *vpn-id*] [**m-vpls**] [**b-vpls** | **i-vpls**] [**etree**] [**name** *name*] [**create**]

no vpls *service-id*

Context config>service

Description This command creates or edits a VPLS instance. If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

A VPLS service connects multiple customer sites together, acting as a zero-hop, Layer 2 switched domain. A VPLS is always a logical full mesh.

When a service is created, the **create** keyword must be specified if the **create** command is enabled in the **environment** context.

When a service is created, the **customer** keyword and *customer-id* must be specified and associated the service with a customer. The *customer-id* must already exist, having been created using the **customer** command in the service context. Once a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and recreated with a new customer association.

To create a management VPLS, the **m-vpls** keyword must be specified.

Once a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified results in an error.

More than one VPLS service may be created for a single customer ID.

By default, no VPLS instances exist until they are explicitly created.

The **no** form of this command deletes the VPLS service instance with the specified *service-id*. The service cannot be deleted until all SAPs and SDPs defined within the service ID have been shutdown and deleted, and the service has been shutdown.

Parameters *service-id* — Specifies the unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every router on which this service is defined.

Values *service-id*: 1 to 214748364

svc-name: A string up to 64 characters.

customer-id — Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn-id — Specifies the VPN ID number which allows you to identify VPNs by a VPN identification number.

Values 1 to 2147483647

Default **null** (0)

m-vpls — Specifies a managed VPLS.

name — Configures an optional service name identifier, up to 64 characters, to a given service. This service name can then be used in configuration references, display, and show commands throughout the system. A defined service name can help the service provider or administrator to identify and manage services within the SR OS platforms.

To create a service, you must assign a service ID; however, after it is created, either the service ID or the service name can be used to identify and reference a service.

If a name is not specified at creation time, then SR OS assigns a string version of the *service-id* as the name.

Service names may not begin with an integer (0 to 9).

create — Keyword used to create the service ID. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

igmp-snooping

Syntax **igmp-snooping**

Context config>service>vpls
config>service>vpls>sap
config>service>vpls>spoke-sdp
config>service>vpls>mesh-sdp

Description This command enables the IGMP snooping context.

mvr

Syntax **mvr**

Context config>service>vpls>igmp-snooping
config>service>vpls>mld-snooping
config>service>vpls>sap>igmp-snooping

Description This command enables the context to configure Multicast VPLS Registration (MVR) parameters.

group-policy

Syntax **group-policy** *policy-name*
no group-policy

Context config>service>vpls>igmp-snp>mvr
config>service>vpls>mld-snp>mvr

Description Identifies filter policy of multicast groups to be applied to this MVR VPLS. The sources of the multicast traffic must be a member of the MVR VPLS.

The **no** form of the command removes the MVR policy association from the VPLS.

Parameters *policy-name* — Specifies the routing policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes. Routing policies are configured in the **config>router>policy-options** context The router policy must be defined before it can be imported.

query-interval

Syntax **query-interval** *seconds*
no query-interval

Context config>service>vpls>igmp-snooping
config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>mld-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping
config>service>vpls>mesh-sdp>mld-snooping

Description If the **send-queries** command is enabled, this command specifies the interval between two consecutive general queries sent by the system on this SAP or SDP.

The configured query interval must be greater than the configured query response interval.

If **send-queries** is not enabled on this SAP or SDP, the configured query interval value is ignored.

The **no** form of the command reverts to the default value.

Default query-interval 125

Parameters *seconds* — Specifies the time interval, in seconds, that the router transmits general host-query messages.

Values 2 to 1024

query-response-interval

Syntax **query-response-interval** *seconds*

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping
config>service>vpls>mesh-sdp>mld-snooping

Description If the **send-queries** command is enabled, this parameter specifies the maximum response time advertised in IGMPv2 or IGMPv3 queries.

The configured query response interval must be smaller than the configured query interval.

If **send-queries** is not enabled on this SAP or SDP, the configured query response interval value is ignored.

The **no** form of the command reverts to the default value.

Default query-response-interval 10

Parameters *seconds* — Specifies the length of time to wait to receive a response to the host-query message from the host.

Values 1 to 1023

report-src-ip

Syntax **report-src-ip** *ip-address*
no report-src-ip

Context config>service>vpls>igmp-snooping

Description This parameters specifies the source IP address used when generating IGMP reports. According the IGMPv3 standard, a zero source address is allowed in sending IGMP reports. However, for interoperability with some multicast routers, the source IP address of IGMP group reports can be configured using this command.

The **no** form of the command reverts to the default value.

Default report-src-ip 0.0.0.0

Parameters *ip-address* — Specifies the source IP source address in transmitted IGMP reports.

query-src-ip

Syntax **query-src-ip** *ipv6-address*

no query-src-ip

Context config>service>vpls>mld-snooping
config>service>vpls>igmp-snooping

Description This command configures the IP source address used in IGMP or MLD queries.

The **no** form of the command removes the IPv6 address from this configuration.

Parameters *ipv6-address* — Specifies an IPv6 address in the form:

x:x:x:x:x:x:x:x:x:x:x:x:d.d.d.d
where:
x - [0 to FF]
d - [0 to 255]

robust-count

Syntax **robust-count** *robust-count*
no robust-count

Context config>service>vpls>igmp-snooping
config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping
config>service>vpls>mesh-sdp>mld-snooping

Description If the **send-queries** command is enabled, this parameter allows tuning for the expected packet loss on a SAP or SDP and is comparable to a retry count. If this SAP or SDP is expected to experience high packet loss, this parameter may be increased. IGMP snooping on this SAP or SDP is robust to packet losses.

If **send-queries** is not enabled, this parameter is ignored.

The **no** form of the command reverts to the default value.

Default robust-count 2

Parameters *robust-count* — Specifies the robust count for the SAP or SDP.

Values 2 to 7

fdb-table-high-wmark

Syntax [**no**] **fdb-table-high-wmark** *high-water-mark*

Context config>service>vpls

Description This command specifies the value to send logs and traps when the threshold is reached.

The **no** form of the command reverts to the default.

Default fdb-table-high-wmark 95

Parameters *high-water-mark* — Specifies the value when to send logs and traps when the threshold is reached.

Values 0 to 100

fdb-table-low-wmark

Syntax [**no**] **fdb-table-low-wmark** *low-water-mark*

Context config>service>vpls

Description This command specifies the value to send logs and traps when the threshold is reached.

The **no** form of the command reverts to the default.

Default fdb-table-low-wmark 90

Parameters *low-water-mark* — Specifies the value to send logs and traps when the threshold is reached.

Values 0 to 100

fdb-table-size

Syntax **fdb-table-size** *table-size*
no fdb-table-size [*table-size*]

Context config>service>vpls

Description This command specifies the maximum number of MAC entries in the forwarding database (FDB) for both learned and static MAC addresses for the VPLS instance.

The **no** form of this command returns the maximum FDB table size to its default.

Default fdb-table-size 250

Parameters *table-size* — Specifies the number of entries permitted in the forwarding database for this VPLS instance.

Values 7750 OS: SR-7, SR-12, SR-12e: 511999
SR-e/SR-a: 250000

mfib-table-high-wmark

Syntax [**no**] **mfib-table-high-wmark** *high-water-mark*

Context config>service>vpls

Description This command specifies the multicast FIB high watermark. When the percentage filling level of the multicast FIB exceeds the configured value, a trap is generated and a log entry is added.

The **no** form of the command reverts to the default.

Default mfib-table-high-wmark 95

Parameters *high-water-mark* — Specifies the multicast FIB high watermark as a percentage.

Values 1 to 100

mfib-table-low-wmark

Syntax [**no**] **mfib-table-low-wmark** *low-water-mark*

Context config>service>vpls

Description This command specifies the multicast FIB low watermark. When the percentage filling level of the multicast FIB drops below the configured value, the corresponding trap is cleared and a log entry is added.

The **no** form of the command reverts to the default.

Default mfib-table-low-wmark 90

Parameters *low-water-mark* — Specifies the multicast FIB low watermark as a percentage.

Values 1 to 100

mfib-table-size

Syntax **mfib-table-size** *size*
no mfib-table-size

Context config>service>vpls

Description This command specifies the maximum number of (s,g) entries in the multicast forwarding database (MFIB) for this VPLS instance.

When a table size limit is set on the MFIB of a service which is lower than the current number of dynamic entries present in the MFIB then the number of entries remains above the limit.

The **no** form of this command removes the configured maximum MFIB table size.

Parameters *size* — Specifies the maximum number of (s,g) entries allowed in the MFIB.

Values 1 to 16383

fast-leave

Syntax [**no**] **fast-leave**

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping
config>service>vpls>mesh-sdp>mld-snooping

Description This command enables fast leave.

When IGMP fast leave processing is enabled, the 7450 ESS or 7750 SR immediately removes a SAP or SDP from the IP multicast group when it detects an IGMP leave message on that SAP or SDP. Fast leave processing allows the switch to remove a SAP or SDP that sends a leave message from the forwarding table without first sending out group-specific queries to the SAP or SDP, which speeds up the process of changing channels.

Fast leave should only be enabled when there is a single receiver present on the SAP or SDP.

When fast leave is enabled, the configured **last-member-query-interval** value is ignored.

import

Syntax **import** *policy-name*
no import

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping
config>service>vpls>mesh-sdp>mld-snooping

Description This command specifies the import routing policy to be used for IGMP packets to be used on this SAP or SDP. Only a single policy can be imported on a SAP at any time.

The **no** form of the command removes the policy association from the SAP or SDP.

Parameters *policy-name* — Specifies the routing policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes. Routing policies are configured in the **config>router>policy-options** context The router policy must be defined before it can be imported.

last-member-query-interval

Syntax **last-member-query-interval** *tenths-of-seconds*
no last-member-query-interval

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping
config>service>vpls>mesh-sdp>mld-snooping

Description This command configures the maximum response time used in group-specific queries sent in response to 'leave' messages, and is also the amount of time between two consecutive group-specific queries. This value may be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group.

The configured **last-member-query-interval** is ignored when fast leave is enabled on the SAP or SDP.

The **no** form of the command reverts to the default value.

Default last-member-query-interval 10

Parameters *seconds* — Specifies the frequency, in tenths of a second, at which query messages are sent.

Values 1 to 50

max-num-groups

Syntax **max-num-groups** *count*
no max-num-groups

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping
config>service>vpls>mesh-sdp>mld-snooping

Description This command defines the maximum number of multicast groups that can be joined on this SAP or SDP. If the 7450 ESS or 7750 SR receives an IGMP join message that would exceed the configured number of groups, the request is ignored.

The **no** form of the command reverts to the default value.

Parameters *count* — Specifies the maximum number of groups that can be joined on this SAP or SDP.

Values 1 to 1000

mrouter-port

Syntax [**no**] **mrouter-port**

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping

Description This command specifies whether a multicast router is attached behind this SAP or SDP.

Configuring a SAP as an mrouter-port has two effects. First, all multicast traffic received on another SAP or SDP is copied to this SAP or SDP. Second, IGMP reports generated by the system as a result of joining or leaving a multicast group, are sent to this SAP or SDP.

If two multicast routers exist in the network, one of them becomes the active querier. While the other multicast router (non-querier) stops sending IGMP queries, it should still receive reports to keep its multicast trees up to date. To support this, **mrouter-port** should be enabled on all SAPs or SDPs connecting to a multicast router.

Note: The IGMP version to be used for the reports (v1, v2 or v3) can only be determined after an initial query has been received. Until this query is received, no reports are sent on the SAP or spoke SDP, even if **mrouter-port** is enabled.

If the **send-queries** command is enabled on this SAP or spoke SDP, the **mrouter-port** parameter can not be set.

The **no** form of the command reverts to the default.

from-vpls

Syntax **from-vpls** *vpls-id*
no from-vpls

Context config>service>vpls>sap>igmp-snp>mvr
config>service>vpls>sap>mld-snp>mvr

Description This command configures the VPLS from which multicast traffic is copied upon receipt of an IGMP join request.

IGMP snooping must be enabled on the MVR VPLS.

Parameters *vpls-id* — Specifies the MVR VPLS from which multicast channels should be copied into this SAP.

to-sap

Syntax **to-sap** *sap-id*
no to-sap

Context config>service>vpls>sap>snooping>mvr

Description In some situations, the multicast traffic should not be copied from the MVR VPLS to the SAP on which the IGMP message was received (standard MVR behavior) but to another SAP.

This command configures the SAP to which the multicast data needs to be copied.

The **no** form of the command reverts to the default value.

Parameters *sap-id* — Specifies the SAP to which multicast channels should be copied.

send-queries

Syntax [**no**] **send-queries**

Context config>service>vpls>sap>igmp-snooping
config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping
config>service>vpls>mesh-sdp>mld-snooping

Description This command specifies whether to send IGMP general query messages on the SAP or SDP.

If **mrouter-port** is enabled on this SAP or spoke SDP, the **send-queries** command parameter can not be set.

The **no** form of the command disables the IGMP general query messages.

static

Syntax **static**

Context config>service>vpls>sap>igmp-snooping

config>service>vpls>spoke-sdp>igmp-snooping
config>service>vpls>mesh-sdp>igmp-snooping
config>service>vpls>sap>mld-snooping
config>service>vpls>spoke-sdp>mld-snooping
config>service>vpls>mesh-sdp>mld-snooping

Description This command enables the context to configure static group addresses. Static group addresses can be configured on a SAP or SDP. When present, either as a (*, g) or a (s,g) entry, multicast packets matching the configuration are forwarded even if no join message was registered for the specific group.

group

Syntax [**no**] **group** *grp-address*

Context config>service>vpls>sap>igmp-snp>static
config>service>vpls>spoke-sdp>igmp-snp>static
config>service>vpls>mesh-sdp>igmp-snp>static

Description This command adds a static multicast group either as a (*, g) or as one or more (s,g) records. When a static IGMP group is added, multicast data for that (*,g) or (s,g) is forwarded to the specific SAP or SDP without receiving any membership report from a host.

Parameters *grp-address* — Specifies an IGMP multicast group address that receives data on an interface. The IP address must be unique for each static group.

source

Syntax [**no**] **source** *ip-address*

Context config>service>vpls>sap>igmp-snp>static>group
config>service>vpls>spoke-sdp>igmp-snp>static>group
config>service>vpls>mesh-sdp>igmp-snp>static>group

Description This command adds a static (s,g) entry to allow multicast traffic for the corresponding multicast group from that specific source. For the same multicast group, more than one source can be specified.

Static (s,g) entries can not be entered when a starg is already created.

The **no** form of the command removes the source IP address from this configuration.

Parameters *ip-address* — Specifies the IPv4 unicast address.

starg

Syntax [**no**] **starg**

Context config>service>vpls>sap>igmp-snp>static>group
config>service>vpls>spoke-sdp>igmp-snp>static>group
config>service>vpls>mesh-sdp>igmp-snp>static>group

Description This command adds a static (*,g) entry to allow multicast traffic for the corresponding multicast group from any source. This command can only be enabled if no existing source addresses for this group are specified.

The **no** form of the command reverts to the default.

8.6.2.1.7 Chassis Level Commands

chassis-level

Syntax **chassis-level**

Context config>mcast-management

Description This command enables the context to configure multicast plane bandwidth parameters. The chassis-level CLI node contains the multicast plane replication limit for each switch fabric multicast plane.

The chassis-level node always exists and contains the configuration command to define the total replication rates for primary and secondary associated ingress paths for each switch fabric multicast plane.

mmrp-impm-override

Syntax [**no**] **mmrp-impm-override**

Context config>mcast-mgmt>chassis-level

Description This command enables ingress Multicast Path Management (IMPM) from monitoring PIM and IGMP.

The **no** form of the command disables the IMPM monitoring.

per-mcast-plane-capacity

Syntax [**no**] **per-mcast-plane-capacity**

Context config>mcast-mgmt>chassis-level

Description This command enables the context to configure multicast plane bandwidth parameters. This CLI node contains the configuration of the overall multicast (primary plus secondary) and specific secondary rate limits for each switch fabric multicast plane.

mcast-capacity

Syntax **mcast-capacity** *primary-percentage* **secondary** *secondary-percentage*
no mcast-capacity

Context config>mcast-mgmt>chassis-level>plane-capacity

Description This command configures the primary and secondary multicast plane capacities used when the full complement of possible switch fabrics in the system is not up (at least one possible switch fabric is not provisioned or is down). The rates are defined as a percentage of the total multicast plane capacity which is configured using the **total-capacity** command.

The **no** form of the command reverts to the default values.

Default

	primary-percentage 100.00
	secondary-percentage 90.00
with SFM5:	*primary-percentage* 87.5
	secondary-percentage 87.5

Parameters *primary-percentage* — Specifies the percentage of the total multicast plane capacity to be used for primary multicast planes.

secondary-percentage — Specifies the percentage of the total multicast plane capacity to be used for secondary multicast planes.

Values 0.01 to 100

redundant-mcast-capacity

Syntax **redundant-mcast-capacity** *primary-percentage* **secondary** *secondary-percentage*
no redundant-mcast-capacity

Context config>mcast-mgmt>chassis-level>plane-capacity

Description This command configures the primary and secondary multicast plane capacities used when the full complement of possible switch fabrics in the system are up. The rates are defined as a percentage of the total multicast plane capacity which is configured using the **total-capacity** command.

The **no** form of the command reverts to the default values.

Default

	primary-percentage 100.00
	secondary-percentage 90.00
with SFM5:	*primary-percentage* 87.5
	secondary-percentage 87.5

Parameters *primary-percentage* — Specifies the percentage of the total multicast plane capacity to be used for primary multicast planes.

Values 0.01 to 100

secondary-percentage — Specifies the percentage of the total multicast plane capacity to be used for secondary multicast planes.

Values 0.01 to 100

total-capacity

Syntax **total-capacity** *capacity*
no total-capacity

Context config>mcast-mgmt>chassis-level>plane-capacity

Description This command configures the total multicast plane capacity supported individually by all switch fabric multicast planes.

The **no** form of the command reverts to the default.

Default SR/ESS: 2000

SR12e: dynamic

Parameters *capacity* — Specifies the multicast plane capacity in Mb/s.

Values **2000**, **4000**, **5250**, **8250**, **dynamic** (Specifies that multicast plane capacity is determined based on provisioned line cards and switch fabrics in the chassis.)

round-robin-inactive-records

Syntax [**no**] **round-robin-inactive-records**

Context config>mcast-mgmt>chassis-level

Description This command specifies whether initially inactive multicast records use the IOM default secondary multicast path or not. When enabled, the system redistributes newly populated inactive records among all available IOM multicast paths and multicast switch fabric planes. When disabled, the system continues to set all initially inactive multicast records to use the IOM default secondary multicast path.

The **no** form of the command reverts to the default.

8.6.2.1.8 Multicast Redirection Commands

action

Syntax **action** {**accept** | **next-entry** | **next-policy** | **drop** | **reject**}
no action

Context config>router>policy-options>policy-statement>entry

Description This command creates the context to configure actions to take for routes matching a route policy statement entry.

This command is required and must be entered for the entry to be active.

Any route policy entry without the **action** command is considered incomplete and is inactive.

The **no** form of the command deletes the action context from the entry.

Parameters **accept** — Specifies routes matching the entry match criteria are accepted and propagated.

next-entry — Specifies that the actions specified would be made to the route attributes and then policy evaluation would continue with next policy entry (if any others are specified).

next-policy — Specifies that the actions specified would be made to the route attributes- and then policy evaluation would continue with next route policy (if any others are specified).

drop — Specifies that routes matching the entry match criteria should be rejected. This parameter provides a context for modifying route properties.

reject — Specifies that routes matching the entry match criteria should be rejected. This parameter does not provide a context for modifying route properties.

interface

Syntax **interface** *port:tags*

Context config>router>igmp

config>services>ies
config>services>vprn

Description This command enables the context to configure IGMP (or MLD) interface parameters on the multicast redirect interface.

8.6.2.1.9 Multicast Query Timer Commands

disable-router-alert-check

Syntax [**no**] **disable-router-alert-check**

Context config>router>igmp>group-interface
config>router>mld>group-interface
config>service>vprn>igmp>group-interface
config>service>vprn>mld>group-interface

Description This command disables router alert checking for IGMP/MLD messages received on this interface.

The **no** form of the command enables router alert checking.

query-interval

Syntax **query-interval** *seconds*
no query-interval

Context config>router>igmp>group-interface
config>router>mld>group-interface
config>service>vprn>igmp>group-interface
config>service>vprn>mld>group-interface

Description This command configures the query interval when the group interface is configured with **no sub-hosts-only**. If nothing is configured, by default, the **query-interval** takes the value defined in the **config**>**router**>**igmp** (or **mld**) context or in the **config**>**service**>**vprn**>**igmp** (or **mld**) context.

The **no** form of the command reverts to the default value.

Default query-interval 125

Parameters *seconds* — Specifies the frequency, in seconds, at which the router transmits general host-query messages.

Values 2 to 1024

query-last-listener-interval

Syntax **query-last-listener-interval** *seconds*
no query-last-listener-interval

Context config>router>mld>group-interface
config>service>vprn>mld>group-interface

Description This command configures the frequency at which the querier router sends a group-specific query messages, including the messages sent in response to leave-group messages and is only applicable when the group interface is configured with the **no sub-hosts-only** command. The shorter the interval, the faster the loss of the last listener of a group can be detected. If nothing is configured, by default, the **query-last-listener-interval** takes the value defined in the **config**>**router**>**mld** context or in the **config**>**service**>**vprn**>**mld** context.

The **no** form of the command reverts to the default value.

Default query-last-listener-interval 1

Parameters *seconds* — Specifies the frequency, in seconds, at which query messages are sent.

Values 1 to 1023

query-last-member-interval

Syntax **query-last-member-interval** *seconds*
no query-last-member-interval

Context config>router>igmp>group-interface
config>service>vprn>igmp>group-interface

Description This command configures the frequency at which the querier router sends group-specific query messages, including the messages sent in response to leave-group messages and is only applicable when the group interface is configured with **no sub-hosts-only**. The shorter the interval, the faster the loss of the last member of a group can be detected. If nothing is configured, by default, the **query-last-member-interval** takes the value defined in the **config**>**router**>**igmp** context or in the **config**>**service**>**vprn**>**igmp** context.

The **no** form of the command reverts to the default value.

Default query-last-member-interval 1

Parameters *seconds* — Specifies the frequency, in seconds, at which query messages are sent.

Values 1 to 1024

query-response-interval

Syntax **query-response-interval** *seconds*

no query-response-interval

Context config>router>igmp>group-interface
config>router>mld>group-interface
config>service>vprn>igmp>group-interface
config>service>vprn>mld> group-interface

Description This command configures the query response interval on when the group interface is configured with the **no sub-hosts-only** command. If nothing is configured, by default, the **query-response-interval** takes the value defined in the **config**>**router**>**igmp** (or **mld**) context or in the **config**>**service**>**vprn**>**igmp** (or **mld**) context.

The **no** form of the command reverts to the default value.

Default query-response-interval 10

Parameters *seconds* — Specifies the length of time to wait to receive a host-query message response from the host.

Values 1 to 1023

8.6.2.1.10 Forwarding Plane Commands

fp

Syntax **fp** [*fp-number*]

Context config>card

Description This command enables access to the configuration of the forwarding planes on a card.

Commands can only be configured under **card>fp** if the hardware that the FP resides on (either a card or an XMA) is provisioned. Conversely, all commands under **card>fp** of the corresponding FPs are automatically removed when that hardware is unprovisioned.

Parameters *fp-number* — This parameter is optional following the **fp** command.

Values 1 to 8

Default 1

ingress

Syntax **ingress**

Context config>card>fp

Description This command enables access to the ingress fp CLI context.

mcast-path-management

Syntax **mcast-path-management**

Context config>card>fp>ingress

Description This CLI node contains the forwarding plane settings for ingress multicast path management. Enter the node to configure the bandwidth-policy and the administrative state of ingress multicast path management.

bandwidth-policy

Syntax **bandwidth-policy** *policy-name*
no bandwidth-policy

Context config>card>fp>ingress>mcast-path-management

Description This command is used to explicitly associate a bandwidth policy to a forwarding plane or MDA. The bandwidth policy defines the dynamic rate table and the multicast paths bandwidth and queuing parameters.

If a bandwidth policy is not explicitly associated with a forwarding plane or MDA, the default bandwidth policy is used when ingress multicast path management is enabled.

The **no** form of the command removes an explicit bandwidth policy from a forwarding plane or MDA and restores the default bandwidth policy.

Parameters *policy-name* — The policy-name parameter is required and defines the bandwidth policy that should be associated with the MDA or forwarding plane for ingress multicast path management. If the policy name does not exist, the bandwidth-policy command will fail. The name can be up to 32 characters.

Values Any existing bandwidth policy name.

Default default

8.6.2.2 Show Commands

mcast-management

Syntax **mcast-management**

Context show

Description This command shows multicast path management related information.

bandwidth-policy

Syntax **bandwidth-policy** *policy-name* [**detail**]

Context show>mcast-management

Description This command displays multicast path management bandwidth policy information.

Parameters *policy-name* — Specifies the bandwidth policy name, up to 32 characters.

detail — Displays detailed information.

Output The following output displays an example of multicast management bandwidth policy information.

Sample Output

```
*A:PE-1# show mcast-management bandwidth-policy "default" detail
===============================================================================
Bandwidth Policy Details
===============================================================================
-------------------------------------------------------------------------------
Policy             : default
-------------------------------------------------------------------------------
Admin BW Thd       : 10 kbps            Falling Percent RST: 50
Mcast Pool Total   : 10                 Mcast Pool Resv Cbs: 50
Slope Policy       : default
Primary
Cbs                : 5.00               Mbs                : 7.00
Low Drop Tail      : 10
Secondary
Cbs                : 30.00              Mbs                : 40.00
Low Drop Tail      : 10                 Paths(Single/Dual) : 1/1
===============================================================================
Bandwidth Policies : 1
===============================================================================
*A:PE-1#
```

channel

Syntax **channel** [**router** *router-instance* | **vpls** s*ervice-id* | **service-name** *service-name*] [fp *slot-number*[*/fp-number*]] [**group** *ip-address* [**source** *ip-address*]] [**path** *path-type*] [**detail**]

Context show>mcast-management

Description This command displays multicast path management channel related information.

Parameters *service-id* — Specifies an existing VPLS service ID.

Values 1 to 214748364

service-name — A string up to 64 characters.

slot-number — Displays information for the specified card slot.

fp-number — Specifies an FP.

Values 1 to 8

ip-address — ipv4-address in the form *a.b.c.d*

path-type — Specifies the path type.

Values primary, secondary

Output The following output displays an example of multicast management channel information.

Sample Output

```
*A:PE# show mcast-management channel
===============================================================================
Multicast Channels
===============================================================================
Legend :  D - Dynamic  E - Explicit
===============================================================================
Source Address                          Slot/Fp   Current-Bw  Path      D/E
 Group Address                                     Highest-Bw  Plane
-------------------------------------------------------------------------------
10.0.0.29                               5/1       3930        Primary    D
 233.252.0.1                                       3973        17
===============================================================================
Multicast Channels : 1
===============================================================================
*A:PE# show mcast-management channel detail
===============================================================================
Multicast Channels
===============================================================================
-------------------------------------------------------------------------------
Source Address     : 10.0.0.29
Group Address      : 233.252.0.1
-------------------------------------------------------------------------------
Slot/Complex       : 5/1               Current Bw         : 3918 kbps
Dynamic/Explicit   : Dynamic           Current Path       : Primary
Oper Admin Bw      : 0 kbps            Current Plane      : 17
Ing last highest   : 3973              Preference         : 0
Black-hole rate    : None              Ing sec highest    : 3970
Time remaining     : 20 seconds         Blackhole          : No
===============================================================================
Multicast Channels : 1
===============================================================================
*A:PE#
```

fp

Syntax **fp** [*slot-number*[*/fp-number*]] [**path** *path-type*]

Context show>mcast-management

Description This command displays multicast path management FP-related information.

Parameters *slot-number* — Displays information for the specified card slot.

fp-number — Specifies an FP.

Values 1 to 8

path-type — Specifies the path type.

Values primary, secondary, blackhole

Output The following output displays an example of multicast management MDA information.

Sample Output

```
*A:PE# show mcast-management fp 5/1
===============================================================================
Fp 5/1
===============================================================================
S/F  Bw-policy                       Type           In-use-Bw          Admin
-------------------------------------------------------------------------------
5/1  reduce-paths                    Primary        3964 Kbps          up
     reduce-paths                    Secondary      0 Kbps             up
     reduce-paths                    Blackhole      0 Kbps             up
-------------------------------------------------------------------------------
*A:PE#
```

mcast-reporting-dest

Syntax **mcast-reporting-dest** [*mcast-reporting-dest-name*]

Context show>mcast-management

Description This command displays multicast path management reporting destination information and applies only to the 7750 SR.

8.6.2.3 Debug Commands

mcast-reporting-dest

Syntax [**no**] **mcast-reporting-dest** [*dest-name*]

Context debug>mcast-mgmt

Description This command debugs multicast path management reporting destinations and applies only to the 7750 SR.

igmp

Syntax [**no**] **igmp** [**host** *ip-address*] [**group** *grp-address*]

Context debug>mcast-mgmt>mcast-rprt-dest

Description This command sets mcast reporting dest debug filtering options and applies only to the 7750 SR.

9 Triple Play Enhanced Subscriber Management

9.1 Uniform RADIUS Server Configuration

9.1.1 RADIUS Server Configuration

The following two configuration methods co-exist but are mutually exclusive:

- Uniform RADIUS Server Configuration (Preferred)
- Legacy RADIUS Server Configuration

9.1.1.1 Uniform RADIUS Server Configuration (Preferred)

This configuration method is preferred as it can be re-used amongst multiple applications (Subscriber authentication and accounting, L2TP tunnel accounting, WLAN gateway RADIUS proxy,) and enables additional functionality not available in the legacy configuration method. For example:

- A RADIUS server policy operational state can be controlled by reception of accounting on or off responses.
- Buffering of accounting messages: When all servers in a **radius-server-policy** are unreachable, it is possible to buffer the acct-stop and acct-interim-update messages for up to 25 hours. When a RADIUS server becomes reachable again then the messages in the buffer are retransmitted.
- A configurable hold down time for accounting servers that are marked down and during which no new communication attempts will be made (hold-down-time).
- A configurable maximum number of outstanding RADIUS requests for accounting servers (pending-requests-limit).
- Increased retry and timeout values for unsuccessful RADIUS communication.
- Enhanced RADIUS server statistics
- IPv6 RADIUS server

Note: A RADIUS server is marked down if it detects a number of consecutive timeouts independent of transaction-id or origin of request.

Where a number of consecutive timeouts is defined by the number of retries configured below the radius-server-policy servers.

The default number of retries is 3, meaning 1 initial try and 2 retries.

If, for example, the RADIUS server has “2 timeouts, 1 reply, 1 timeouts”, whereby the timeouts are originated for the same host, the server is not marked down since intermediate replies were received.

To attach a RADIUS server policy to an authentication policy:

For example,

```
configure
    subscriber-mgmt
        authentication-policy "auth-policy-1" create
            radius-server-policy "aaa-server-policy-1"
        exit
    exit
```

Note:

- To avoid conflicts, the following CLI commands are ignored in the authentication policy when a **radius-server-policy** is attached:
  - All commands in the **radius-authentication-server** context
  - **accept-authorization-change**
  - **coa-script-policy**
  - **accept-script-policy**
  - **request-script-policy**
- The **fallback-action** command specifies the action when no RADIUS server is available is configured direct in the **config>subscr-mgmt>auth-plcy** CLI context.

To attach a RADIUS server policy to a RADIUS accounting policy:

For example:

```
configure
    subscriber-mgmt
        radius-accounting-policy "acct-policy-1" create
            radius-server-policy "aaa-server-policy-1"
        exit
```

```
    exit
```

Note: To avoid conflicts, the following CLI commands are ignored in the RADIUS accounting policy when a **radius-server-policy** is attached:

- All commands in the **radius-accounting-server** context
- **acct-request-script-policy**

To configure the RADIUS servers in a RADIUS server policy:

For example:

```
configure
    aaa
        radius-server-policy "aaa-server-policy-1" create
            description "Radius AAA server policy"
            accept-script-policy "script-policy-2"
            acct-on-off oper-state-change
            acct-request-script-policy "script-policy-3"
            auth-request-script-policy "script-policy-1"
            no python-policy
            servers
                access-algorithm direct
                hold-down-time sec 30
                no ipv6-source-address
                retry 3
                router "Base"
                no source-address
                timeout sec 5
                buffering
                    acct-interim min 60 max 3600 lifetime 5
                    acct-stop min 60 max 3600 lifetime 5
                exit
                server 1 name "server-1"
                server 2 name “server-2”
            exit
        exit
    exit
```

To configure the RADIUS servers in the routing instance:

- In the Base routing instance: **config>router>radius-server**.
- In a VPRN routing instance: **config>service>vprn 10>radius-server**.
- In the management routing instance (out of band): **config>router management>radius-server**.

For example:

```
configure
```

```
router
    radius-server
        server "server-1" address 172.16.1.1 secret <shared secret> hash2 create
            accept-coa
            coa-script-policy "script-policy-4"
            description "Radius server 1"
            pending-requests-limit 4096
            acct-port 1813
            auth-port 1812
        exit
        server "server-2" address 172.16.1.2 secret <shared secret> hash2 create
            accept-coa
            coa-script-policy "script-policy-4"
            description "Radius server 2"
            pending-requests-limit 4096
            acct-port 1813
            auth-port 1812
        exit
    exit
exit
```

Note: To configure RADIUS CoA servers for use in Enhanced Subscriber Management, the server must be configured in the corresponding routing instance with the **accept-coa** command enabled.

9.1.1.2 Legacy RADIUS Server Configuration

Note: It is recommended to migrate to the uniform RADIUS server configuration as described above to have additional functionality enabled.

To configure a RADIUS server in an authentication policy:

```
configure
    subscriber-mgmt
        authentication-policy "auth-policy-1" create
            radius-authentication-server
                access-algorithm direct
                hold-down-time 30
                retry 3
                no source-address
                timeout 5
                router "Base"
                server 1 address 172.16.1.1 secret <shared secret> hash2 port 1812
                    pending-requests-limit 4096
                server 2 address 172.16.1.2 secret <shared secret> hash2 port 1812
                    pending-requests-limit 4096
            exit
```

```
            accept-authorization-change
            accept-script-policy "script-policy-2"
            coa-script-policy "script-policy-4"
            request-script-policy "script-policy-1"
        exit
    exit
```

Note: In legacy RADIUS server configuration, to configure RADIUS CoA servers for use in Enhanced Subscriber Management, the server must be configured in the authentication policy with the **accept-authorization-change** command enabled. A CoA only server can be configured with the optional coa-only flag.

To configure a RADIUS server in a RADIUS accounting policy:

```
configure
    subscriber-mgmt
        radius-accounting-policy "acct-policy-1" create
            radius-accounting-server
                access-algorithm direct
                retry 3
                timeout 5
                no source-address
                router "Base"
                server 1 address 172.16.1.1 secret <shared secret> hash2 port 1813
                server 2 address 172.16.1.2 secret <shared secret> hash2 port 1813
            exit
            acct-request-script-policy "script-policy-3"
        exit
    exit
```

9.2 RADIUS Authentication of Subscriber Sessions

This section describes the Nokia router acting as a Broadband Subscriber Aggregator (BSA).

Note: In the TPSDA solutions, the Nokia 5750 Subscriber Services Controller (SSC) serves as the policy manager, DHCP and RADIUS server.

In this application, one of the required functions can be to authenticate users trying to gain access to the network. While sometimes the DHCP server (an SSC) can perform authentication, in most cases a RADIUS server (an SSC) is used to check the customer's credentials.

Note: See the DHCP Management section for information about DHCP and DHCP Snooping.

For information about the RADIUS server selection algorithm, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR System Management Guide*.

If authentication is enabled, the router will temporarily hold any received DHCP discover message and will send a access-request message to a configured RADIUS server containing the client's MAC address or Circuit-ID (from the Option 82 field) as the user name. If and when access is granted by the RADIUS server, the router then forwards or relays the DHCP discover message to the DHCP server and allows an IP address to be assigned. If the RADIUS authentication request is denied, the DHCP message is dropped and an event is generated.

A typical initial DHCP scenario (after client boot up) is shown in Figure 65.

Figure 65 **Initial DHCP Scenario**

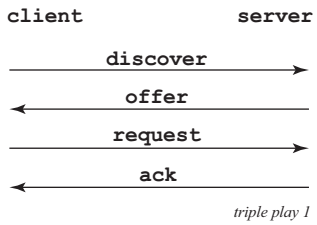

But, when the client already knows its IP address (when an existing lease is being renewed), it can skip straight to the request/ack phase, as shown in Figure 66.

Figure 66 **DHCP Scenario with Known IP Address**

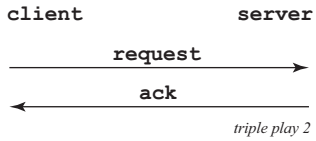

In the first scenario, the DHCP discover triggers an authentication message to RADIUS and the DHCP request also triggers RADIUS authentication. The previous reply is cached for 10 seconds, the second DHCP packet will not result in a RADIUS request.

In the second scenario, the DHCP request triggers an authentication message to RADIUS.

If the optional subscriber management authentication policy **re-authentication** command is enabled, DHCP authentication is performed at every DHCP lease renew request. Only dynamic DHCP sessions are subject to remote authentication. Statically provisioned hosts are not authenticated.

9.2.1 RADIUS Authentication Extensions

This section describes an extension to RADIUS functionality in the subscriber management context. As part of subscriber host authentication, RADIUS can respond with access-response message, which, in the case of an accept, can include several RADIUS attributes (standard and vendor-specific) that allow proper provisioning of a given subscriber-host.

Change-of-Authorization (CoA) messages as defined by RFC 3576, *Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS)*, are supported. The goal of CoA messages is to provide a mechanism for “mid-session change” support through RADIUS.

9.2.1.1 Triple Play Network with RADIUS Authentication

Figure 67 shows a flow of RADIUS authentication of DHCP hosts in the triple play aggregation environment. Besides granting the authentication of given DHCP host, the RADIUS server can include RADIUS attributes (standard and/or Vendor-Specific Attributes (VSAs)) which are then used by the network element to provision objects related to a given DHCP host.

Figure 67 **Triple Play Aggregation Network with RADIUS-Based DHCP Host Authentication**

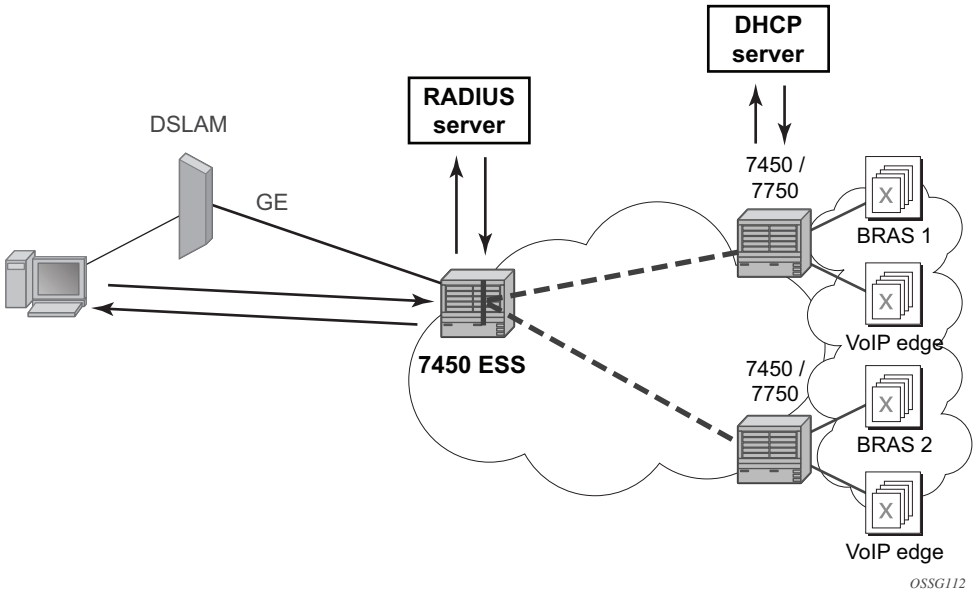

RADIUS is a distributed client/server concept that is used to protect networks against unauthorized access. In the context of the router's subscriber management in TPSDA, the RADIUS client running on nodes sends authentication requests to the SSC.

RADIUS can be used to perform three distinct services:

- Authentication determines whether or not a given subscriber-host is allowed to access a specific service.
- Authorization associates connection attributes or characteristics with a specific subscriber host.
- Accounting tracks service use by individual subscribers.

The RADIUS protocol uses “attributes” to describe specific authentication, authorization, and accounting elements in a user profile (which are stored on the RADIUS server). RADIUS messages contain RADIUS attributes to communicate information between network elements running a RADIUS client and a RADIUS server.

RADIUS divides attributes into two groups, standard attributes and Vendor-Specific Attributes (VSAs). VSA is a concept allowing conveying vendor-specific configuration information in a RADIUS messages, as discussed in RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*. It is up to the vendor to specify the exact format of the VSAs.

Nokia-specific VSAs are identified by vendor-id 6527.

9.2.2 RADIUS Authorization Extensions

The following sections define different functional extensions and list relevant RADIUS attributes.

Basic Provisioning of Authentication Extensions

In order to comply with RFC 4679, *DSL Forum Vendor-Specific RADIUS Attributes*, the software includes the following attributes in the authentication-request message:

- agent-circuit-id (as defined by DSL forum)
- agent-remote-id (as defined by DSL forum)

The following attributes can also be included if configured and provided by downstream equipment:

- actual-data-rate-upstream
- actual-data-rate-downstream
- minimum-data-rate-upstream
- minimum-data-rate-downstream
- access-loop-encapsulation

When the node is configured to insert (or replace) Option 82, the above mentioned attributes will have the content after this operation has been performed by the software.

In addition, the following standard RADIUS attributes will be included in authentication request messages (subject to configuration):

- NAS-identifier — string containing system-name
- NAS-port-id
- NAS-port-type — Values: 32 (null encap), 33 (dot1q), 34 (qinq), 15 (DHCP hosts), specified value (0 — 255)
- MAC-address (Nokia VSA 27)
- dhcp-vendor-class-id (Nokia VSA 36)
- calling-station-id

These will only be included in the access-request if they have been configured.

In order to provide the possibility to push new policies for currently active subscribers, the routers support commands to force re-authentication of the given subscriber-host. After issuing such a command, the router will send a DHCP FORCERENEW packet, which causes the subscriber to renew its lease (provided it supports force-renew). The DHCP request and ACK are then authenticated and processed by the routers as they would be during a normal DHCP renew.

9.2.2.1 Calling-Station-ID

A **calling-station-id** can be configured at SAP level and can be included in the RADIUS authentication and accounting messages. This attribute is used in legacy BRAS to identify the user (typically phone number used for RAS connection). In the broadband networks this was replaced by circuit-id in Option 82. However, the Option 82 format is highly dependent on access-node vendor, which makes interpretation in management servers (such as RADIUS) troublesome. Some operators use the calling-station-id attribute as an attribute indicating the way the circuit-id should be interpreted. The calling-station-id attribute can be configured as a string which will be configured on the SAP. It can also be configured to use the **sap-id**, **remote-id** or **mac-address**.

9.2.2.2 Subscriber Session Timeout

To limit the lifetime of a PPP session or DHCPv4 host to a fixed time interval, a timeout can be specified from RADIUS. By default, a PPP session or DHCPv4 host has no session timeout (infinite).

For PPP sessions, a session-timeout can be configured in the ppp-policy. A RADIUS specified session-timeout overrides the CLI configured value.

```
subscriber-mgmt
    ppp-policy "ppp-policy-1" create
```

```
            session-timeout 86400
        exit
    exit
```

When the session timeout expires a PPP session is terminated and a DHCPv4 host deleted. For a DHCPv4 host, a DHCP release message is also sent to the server.

The following two attributes can be used in RADIUS Access-Accept and CoA messages to limit the PPP session or DHCPv4 host session time (Table 71):

Table 71 **Subscriber Session Timeout**

Attribute ID	Attribute Name	Type	Limits	Purpose and Format
27	Session-Timeout	integer	2147483647 seconds	0 = infinite (no session-timeout) (1 to 2147483647) in seconds For example: Session-Timeout = 3600
26-6527-160	Alc-Relative-Session-Timeout	integer	[0 to 2147483647] seconds	0 = infinite (no session-timeout) (1 to 2147483647) in seconds For example: Alc-Relative-Session-Timeout = 3600

When specified in a RADIUS Access-Accept message, both attributes specify an absolute value for session timeout. When specified in a RADIUS CoA message, attribute [26-6527-160] Alc-Relative-Session-Timeout specifies a relative session timeout value in addition to the current session time while attribute [27] Session-Timeout specifies an absolute session timeout value. If the current session time is greater than the received Session-Timeout, a CoA NAK is sent with error cause “Invalid Attribute Value (407)”.

Only one of the above attributes to specify a session timeout can be present in a single RADIUS message. An event is raised when both are specified in a single message.

The output of the **show service id** *service-id* **ppp session detail** CLI command contains following fields related to session timeout for PPP sessions:

- Up Time: the PPP session uptime
- Session Time Left: the remaining time before the session is terminated
- RADIUS Session-TO: the RADIUS received session timeout value.

The output of the **show service id** *service-id* **dhcp lease-state detail** CLI command contains following fields related to session timeout for DHCPv4 hosts:

- Up Time: the DHCPv4 host uptime
- Remaining Lease Time: the remaining time before the lease expires in the DHCP server. The client should renew its lease before this time.
- Remaining SessionTime: the remaining time before the DHCPv4 host is deleted
- Session-Timeout: the DHCPv4 host is deleted when its uptime reaches the Session-Timeout value.
- Lease-Time: the lease time specified by the DHCPv4 server

Note: In a radius-proxy scenario or when a DHCPv4 host is created with a RADIUS CoA message, the RADIUS attribute [26-6527-174] Alc-Lease-Time attribute must be used to specify the lease time. If the [26-6527-174] Alc-Lease-Time is not present in these scenarios, then the RADIUS attribute [27] Session-Timeout is interpreted as DHCPv4 lease time.

9.2.2.2.1 Domain Name in Authentication

In many networks, the user name has specific meaning with respect to the domain (ISP) where the user should be authenticated. In order to identify the user correctly, the user name in an authentication-request message should contain a domain-name. The domain-name can be derived from different places. In PPPoE authentication the domain name is given by the PPPoE client with the user name used in PAP or CHAP authentication. For DHCP hosts similar functionality is implemented by a “pre-authentication” lookup in a local user database before performing the RADIUS request.

For example, it can be derived from option60 which contains the vendor-specific string identifying the ISP the set-box has been commissioned by.

To append a domain name to a DHCP host, the following configuration steps should be taken:

- Under the (group or IP) interface of the service, a local user database should be configured in the DHCP node and no authentication policy should be configured.
- In the local user database, there should be a host entry containing both the domain name to be appended and an authentication policy that should be used for RADIUS authentication of the host. The host entry should contain no other information needed for setting up the host (IP address, ESM string), otherwise the DHCP request is dropped.

- In the authentication policy, the **user-name-format** command should contain the parameter **append** *domain-name*.

9.2.2.3 RADIUS Reply Message for PPPoE PAP/CHAP

The string returned in a [18] Reply-Message attribute in a RADIUS Access-Accept is passed to the PPPoE client in the CHAP Success or PAP Authentication-Ack message.

The string returned in a [18] Reply-Message attribute in a RADIUS Access-Reject is passed to the PPPoE client in the CHAP Failure or PAP Authentication-Nak message.

When no [18] Reply-Message attribute is available, the SR OS default messages are used instead: “CHAP authentication success” or “CHAP authentication failure” for CHAP and “Login ok” or “Login incorrect” for PAP.

9.2.2.4 SHCV Policy

SHCV policies are used to control “subscriber host connectivity verification” which verifies the host connectivity to the BNG. There are two types of SHCV: periodic and event triggered. Prior to Release 13.0R4, some event triggered SHCV relied on the reference timer set by the host-connectivity-verify under the group interface while others had hard-coded values. Release 13.0R4 introduced the “SHCV policy” that allows individual configuration of trigger SHCV timers and periodic SHCV timers depending on the application.

9.2.3 radius-server-policy Retry Attempt Overview

This feature maximizes the use of the remaining healthy RADIUS servers for subscriber authentication and accounting. After the hold-down time expires, a single RADIUS message is used to determine the status of the RADIUS server. If the server remains unresponsive after waiting for a single timeout interval (without any retries), then it is placed back into the hold-down state. If the RADIUS server responds, then it is used for subscriber authentication and accounting with the rest of the healthy servers.

9.2.4 AAA RADIUS Server Operation Status

The different operating states of a RADIUS server are shown in RADIUS Server Operating States. When a RADIUS server is first provisioned into the AAA using the **radius-server-policy** command, the operating state is “unknown”. This state indicates that the RADIUS server has yet to receive a RADIUS request message. To send a request message, the **radius-server-policy** command provides three different access algorithms: direct, round-robin, and hash. With the direct algorithm, request messages are always sent to the in-service RADIUS server with the lowest configured server index. With the round-robin algorithm, the RADIUS requests are load-balanced in a round-robin manner. The hash algorithm offers a load-balanced alternative; the 7750 SR generates a hash-key based on the subscriber information, and the RADIUS request is then sent to a particular server based on the hash key. The hash method differs from the round-robin method in that, under normal working conditions, RADIUS requests from a particular subscriber are always forwarded to the same RADIUS server. When a server replies to a RADIUS request, it transitions from the operational state of “unknown” to “in-service”. A server may transition from “unknown” to “out-of-service” if the server fails to respond to the initial RADIUS message.

Figure 68 **RADIUS Server Operating States**

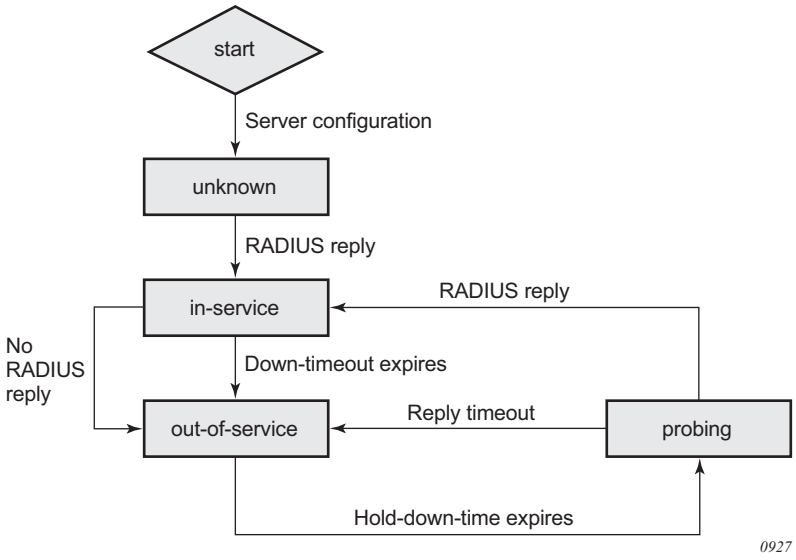

A RADIUS server is declared “out-of-service” when the **down-timeout** timer expires. The router starts the **down-timeout** timer when an access-request is sent. The timer only resets to “0” when a reply is received from the RADIUS server. This means that the timer can be reset to “0” if a reply message is received for another subscriber. For example, the RADIUS server may miss a message but stay “in-service” as long as the server responds to another access request from a different subscriber or from a retry of the same subscriber, as long as the reply is received within the **down-timeout** interval.

Note: It is highly recommended that the **down-timeout** command be set to its default value, which is **no down-timeout**.

The **down-timeout** default value is the timeout value multiplied by the number of retry attempts. The timeout value is the time that the router waits for the RADIUS server to reply, and the retry value is the number of attempts the 7750 SR will make to contact the RADIUS server after each timeout. If the RADIUS server remains unresponsive, the timer will continue to increment until it reaches the configured **down-timeout** value and the server is declared “out-of-service”.

For RADIUS servers that do not respond to all RADIUS requests, a test user account can be optionally set up to periodically send RADIUS request messages in order to keep the server in service. Typically, a RADIUS server should always respond to all access requests. However, creating a test user account for periodic keep-alive may place an unnecessary load on the processor and may lower the overall scale of the router.

At the start of the out-of-service state, a **hold-down-time** timer starts. The timer holds down the RADIUS server and prevents it from operating; no RADIUS messages are sent to an out-of-service server. This is beneficial for the following reasons.

- The server may be unresponsive due to excessive RADIUS message requests; holding it down allows the server to recover.
- Holding down an unresponsive server allows other healthy RADIUS servers to service new requests promptly.

After the **hold-down-time** timer expires, the server enters into the “probing” state. There must be multiple RADIUS servers and at least one healthy server in order for the server to enter the probing state. Probing is always performed by the test user account; actual subscriber requests are never used during probing. If no test user account exists, an actual subscriber request is used to perform the probe. There are no retry attempts; only a single RADIUS message is used to probe a RADIUS server. If the RADIUS server responds, it is declared “in-service” immediately. If the RADIUS

server fails to respond within the timeout value, it is declared “out-of-service” again and the **hold-down-time** timer restarts. Subscriber RADIUS messages used for probing are not cached, and if the server fails to respond, the subscriber is required to send the RADIUS message again by sending an address request; for example, DHCP, PPP, or SLAAC or by performing a **data-trigger**.

9.2.5 AAA RADIUS Accounting Server Stickiness

Stickiness applies to the following subscriber RADIUS accounting sessions: start, interim, and stop. By default, the subscriber sticks with the server that served its last accounting message. For example, if server 1 served the subscriber an accounting start message, then the subsequent interim messages and stop message from the same subscriber will be sent to server 1. If server 1 is out of service, server 2 is used for the subsequent interim and stop messages. When server 1 recovers, the interim and stop messages will stick with server 2. The RADIUS accounting messages will always be forwarded to the server that serviced the subscriber’s last accounting message.

Typically, when using the direct access algorithm, the primary server (lowest configured server index) serves all RADIUS request messages. The other RADIUS servers are used for backup purposes only and might be using a lighter-weight processor. Therefore, it is best to revert to the primary server as soon as it is restored. This can be accomplished by disabling stickiness in direct mode; the RADIUS accounting messages will be forwarded to the primary server once it is restored.

In a round-robin algorithm, while each subscriber session is assigned to a different server in round-robin order, a particular subscriber sticks with a particular server for the entire accounting session. Disabling stickiness will send a subscriber’s RADIUS accounting messages to the list of configured RADIUS servers in a round-robin order.

9.2.6 AAA RADIUS Authentication Fallback Action

The fallback action comes into effect when connectivity to all RADIUS servers is lost. The operating state of the RADIUS servers changes to either “out-of-service” or “probing”. There are two configurable fallback actions: **accept** or **user-db**. An accept action without force-probing will automatically accept all authentication requests from all subscribers. A user-db action without force-probing will use the **local-user-db** for subscriber authentication.

Both **accept** and **user-db** can be combined with the **force-probing** command. Force-probing forces the out-of-service server to transition to the probing state immediately, bypassing the **hold-down-time** timer. Force-probing is a mechanism to promptly restore connectivity to a RADIUS server. A test user is not used to perform a force probe; only actual subscriber authentication will be used to test the operating state of the RADIUS server. Probing only occurs when a server is out of service. If all servers are in the probing state, all new incoming authentication requests follow the fallback action immediately.

When probing with an actual subscriber authentication, the 7750 SR only waits for a reply for one timeout interval without any retries. During the wait, the server is in a probing state and no other subscribers will be used to probe this server. The subscriber authentication request is not cached when used for probing. Therefore, to trigger authentication again, the subscriber is required to authenticate again with an address request or a data-trigger packet.

9.2.7 AAA Test User Account

A test user account is used in the rare case where a RADIUS server ignores RADIUS messages as mentioned in the AAA RADIUS Server Operation Status section. Consequently, when messages are ignored, the router places the RADIUS server out of service. The test user account can keep a RADIUS server in service by periodically sending RADIUS requests to the server. The RADIUS server, while randomly ignoring other subscriber RADIUS requests, must respond to the test user requests. A RADIUS server is in service if it replies to RADIUS messages before the **down-timeout** timer expires. The default **down-timeout** default value is the timeout value multiplied by the retry value, but it is also configurable. The test user account has a configurable interval value, and it is recommended that this value be configured to be less than the **down-timeout** value for it to be useful. The test user account only applies to RADIUS authentication.

Typically, a RADIUS server always responds to all RADIUS requests, and therefore it is not recommended that a test user account be used unless it is absolutely required for certain types of servers. The test user account creates extra load for the processor and can affect scaling. The test user account can be used with a python script (for example, adding additional attributes to the test user account during an **access-request** operation).

9.2.8 Troubleshooting the RADIUS Server

The **tools**>**perform**>**security**>**authentication-server-check** command can be used to troubleshoot a RADIUS server by checking the connectivity and functional status of a RADIUS server for subscriber management operations. The command keyword **debug** can be specified to view additional information on the access request. All VSAs sent and received from the RADIUS server, the hex dump, and all other debug detailed information can be shown without the need to turn on system-wide debugging.

Additional attributes in an Access-Request message can be specified in an attribute file referenced with the command keyword **attr-from-file** *file-url*. Each attribute must be specified on a separate line in the text file in the following format shown in Table 72.

Table 72 **authentication-server-check Attribute File Format**

Attribute file format	Description
<type> = <value>	Standard attribute
<vendor>,<type> = <value>	Vendor Specific Attribute
e,<type>,<ext-type> = <value>	Extended type attribute (RFC 6929)
evs, <type>, <vendor>, <vendortype> = <value>	Extended Vendor Specific attribute (RFC 6929)
le,<type>,<ext-type> = <value>	Long Extended type attribute (RFC6929)
evs, <type>, <vendor>, <vendortype> = <value>	Long Extended Vendor Specific attribute (RFC 6929)

9.2.9 Provisioning of Enhanced Subscriber Management (ESM) Objects

In the ESM concept on network elements, a subscriber host is described by the following aspects:

- subscriber-id-string
- subscriber-profile-string
- sla-profile-string
- ancp-string

- intermediate-destination-identifier-string
- application-profile-string

This information is typically extracted from DHCP-ACK message using a Python script, and is used to provision subscriber-specific resources such as queues and filter entries. As an alternative to extracting this information from DHCP-ACK packet, provisioning from RADIUS server is supported.

As a part of this feature, the following VSAs have been defined:

- alc-subscriber-id-string — Contains a string which is interpreted as a subscriber-id.
- alc-subscriber-profile-string — Contains a string which is interpreted as a subscriber profile
- alc-sla-profile-string — Contains string which is interpreted as an SLA profile.
- alc-ancp-string — Contains string which is interpreted as an ANCP string.
- alc-int-dest-id-string — Contains a string which is interpreted as an intermediate destination ID
- alc-app-profile-string — Contains a string which is interpreted as an application profile

  Note that these strings can be changed in a CoA request.

When RADIUS authentication response messages contain the above VSAs, the information is used during processing of DHCP-ACK message as an input for the configuration of subscriber-host parameters, such as QoS and filter entries.

If ESM is not enabled on a given SAP, information in the VSAs is ignored.

If ESM is enabled and the RADIUS response does not include all ESM-related VSAs (an ANCP string is not considered as a part of ESM attributes), only the subscriber-id is mandatory (the other ESM-related VSAs are not included). The remaining ESM information (sub-profile, sla-profile) will be extracted from DHCP-ACK message according to normal flow (Python script, and so on).

If the profiles are missing from RADIUS, they are not extracted from the DHCP data with Python to prevent inconsistent information. Instead, he data will revert to the configured default values.

However, if the above case, a missing subscriber ID will cause the DHCP request to be dropped. The DHCP server will not be queried in that case.

When no DHCP server is configured, DHCP-discover/request messages are discarded.

9.2.9.1 Provisioning IP Configuration of the Host

The other aspect of subscriber-host authorization is providing IP configuration (ip-address, subnet-mask, default gateway and dns) through RADIUS directory rather than using centralized DHCP server. In this case, the node receiving following RADIUS attributes will assume role of DHCP server in conversation with the client and provide the IP configuration received from RADIUS server.

These attributes will be accepted only if the system is explicitly configured to perform DHCP-server functionality on a given interface.

The following RADIUS attributes will be accepted from authentication-response messages:

- framed-ip-address — The IP address to be configured for the subscriber-host
- framed-ip-netmask — The IP network to be configured for the subscriber host If RADIUS does not return a netmask, the DHCP request is dropped
- framed-pool — The pool on a local DHCP server from which a DHCP-provided IP address should be selected
- alc-default-router — The address of the default gateway to be configured on the DHCP client
- alc-primary-dns — The DNS address to be provided in DHCP configuration.
  - Juniper VSA for primary DNS
  - Redback VSA for primary DNS
- alc-secondary-dns
  - Juniper VSA for secondary DNS
  - Redback VSA for secondary DNS
- alc-lease-time — Defines the lease time
- session-timeout — Defines the lease time in absence of the alc-lease-time attribute
- NetBIOS
  - alc-primary-nbns
  - alc-secondary-nbns

9.2.9.2 RADIUS-Based Authentication in Wholesale Environment

In order to support VRF selection, the following attributes are supported:

- alc-retail-serv-id — Indicates the service-id of the required retail VPRN service configured on the system

9.2.9.3 Change of Authorization and Disconnect-Request

In a typical RADIUS environment, the network element serves as a RADIUS client, which means the messages are originated by a routers. In some cases, such as mid-session changes, it is desirable that the RADIUS server initiates a CoA request to impose a change in policies applicable to the subscriber, as defined by RFC 3576.

To configure a RADIUS server to accept CoA and Disconnect Messages is achieved in one of the following ways:

1. Configure up to 64 RADIUS CoA servers per routing instance:

```
config>router>radius-server#
config>service>vprn>radius-server#
 server "coa-1" address 10.1.1.1 secret <shared-secret> hash2 create
     accept-coa
 exit
```

   This is the preferred method.

2. Configure up to 16 RADIUS CoA servers per authentication policy.

```
config>subscr-mgmt>auth-plcy#
 accept-authorization-change
```

   The UDP port for CoA and Disconnect Messages is configurable per system with the command:

```
config>aaa#
radius-coa-port {1647|1700|1812|3799}
```

Note: There is a priority in the functions that can be performed by CoA. The first matching one will be performed:

- If the CoA packet contains a Force-Renew attribute, the subscriber gets a FORCERENEW DHCP packet. This function is not supported for PPPoE or ARP hosts.
- If the CoA packet contains a create-host attribute, a new lease-state is created. Only DHCP lease-states can be created by a CoA message. PPPoE sessions and ARP hosts cannot be created.
- Otherwise, the ESM strings are updated.

There are several reasons for using RADIUS initiated CoA messages:

1. Changing ESM attributes (SLA or subscriber profiles) or queues/policers/schedulers rates of the given subscriber host — CoA messages containing the identification of the given subscriber-host along with new ESM attributes.
2. Changing (or triggering the change) of IP configuration of the given subscriber-host — CoA messages containing the identification of the given subscriber-host along with VSA indicating request of FORCERENEW generation.
3. Configuring new subscriber-host — CoA messages containing the full configuration for the given host.

If the changes to ESM attributes are required, the RADIUS server sends CoA messages to the network element requesting the change in attributes included in the CoA request:

- attribute(s) to identify a single or multiple subscriber host(s): NAS-Port-Id + IP address/prefix or Acct-Session-Id or Alc-Subsc-ID-Str
  - Nas-Port-Id attribute + single IP address/prefix attribute:
    - Framed-IP-Address
    - Alc-Ipv6-Address
    - Framed-Ipv6-Prefix
    - Delegated-Ipv6-Prefix
    - Alc-Client-Hardward-Addr (Required for private-retail subnet. For more information, refer to the *RADIUS Attributes Reference Guide*.)
  - Acct-Session-Id (number format)
  - Alc-Subsc-ID-Str
- alc-subscriber-profile-string
- alc-sla-profile-string
- alc-ancp-string
- alc-app-profile-string
- alc-int-dest-id-string
- alc-subscriber-id-string
- alc-subscriber-qos-override

Note: If the subscriber-id-string is changed while the ANCP string is explicitly set, the ANCP-string **must** be changed simultaneously. When changing the alc-subscriber-id-string, the lease state is temporarily duplicated, causing two identical ANCP-strings to be in the system at the same time. This is not allowed.

As a reaction to such message, the router changes the ESM settings applicable to the given host.

If changes to the IP configuration (including the VRF-id in the case of wholesaling) of the given host are needed, the RADIUS server may send a CoA message containing VSA indicating request for FORCERENEW generation:

- attribute(s) to identify a single or multiple subscriber host(s): “NAS-Port-Id + IP address/prefix” or “Acct-Session-Id” or “Alc-Subsc-ID-Str”:
  - Nas-Port-Id attribute + single IP address/prefix attribute:
    - Framed-IP-Address
    - Alc-Ipv6-Address
    - Framed-Ipv6-Prefix
    - Delegated-Ipv6-Prefix
  - Acct-Session-Id (number format)
  - Alc-Subsc-ID-Str
- alc-force-renew
- alc-force-nak

As a reaction to such message, router will generate a DHCP FORCERENEW message for the given subscriber host. Consequently, during the re-authentication, new configuration parameters can be populated based on attributes included in Authentication-response message. The force-NAK attribute has the same function as the Force-Renew attribute, but will cause the ESR to reply with a NAK to the next DHCP renew. This will invalidate the lease state on the ESR and force the client to completely recreate its lease, making it possible to update parameters that cannot be updated through normal CoA messages, such as IP address or address pool.

If the configuration of the new subscriber-host is required, RADIUS server will send a CoA message containing VSA request new host generation along with VSAs specifying all required parameters.

- alc-create-host
- alc-subscriber-id-string — This attribute is mandatory in case ESM is enabled, and optional for new subscriber host creation otherwise.
- NAS-port-id — This attribute indicates the SAP where the host should be created.
- framed-ip-address — The framed IP address.
- alc-client-hw-address — A string in the xx:xx:xx:xx:xx:xx format. This attribute is mandatory for new subscriber-host creation.
- alc-lease-time — Specifies the lease time. If both session-timeout and alc-lease-time are not present, then a default lease time of 7 days is used.

- session-timeout — Specifies the lease time in absence of the alc-lease-time attribute. If both session-timeout and alc-lease-time are not present, then a default lease time of 7 days is used.
- alc-retail-svc-id — This is only used in case of wholesaling for selection of the retail service. Indicates the service-id of the required retail VPRN service configured on the system.
- Optionally other VSAs describing given subscriber host. Obviously, if the ESM is enabled, but the CoA message does not contain ESM attributes the new host will not be created.

After executing the requested action, the router element responds with an ACK or NAK message depending on the success/failure of the operation. In case of failure (and hence NAK response), the element will include the error code in accordance with RFC 3576 definitions if an appropriate error code is available.

Supporting CoA messages has security risks as it essentially requires action to unsolicited messages from the RADIUS server. This can be primarily the case in an environment where RADIUS servers from multiple ISPs share the same aggregation network. To minimize the security risks, the following rules apply:

- Support of CoA messages is disabled by default. They can be enabled on a per RADIUS server or authentication-policy basis.
- When CoA is enabled, the node will listen and react only to CoA messages received from RADIUS servers. In addition, CoA messages must be protected with the key corresponding to the given RADIUS server. All other CoA messages will be silently discarded.

In all cases (creation, modification, force-renew) subscriber host identification attributes are mandatory in the CoA request: “NAS-Port-Id + IP” or “Acct-Session-Id” or “Alc-Subsc-ID-Str”

- Nas-Port-Id + single IP address/prefix:
  - Nas-Port-Id
  - Framed-IP-Address
  - Alc-Ipv6-Address
  - Framed-Ipv6-Prefix
  - Delegated-Ipv6-Prefix
  - Alc-Client-Hardware-Addr (may be required for private retail subnet. For more information, refer to the *RADIUS Attributes Reference Guide*.)
- Acct-Session-Id (number format)
- Alc-Subsc-ID-Str

When there are no subscriber host identification attributes present in the CoA, the message will be NAK'd with corresponding error code.

- hosts, meaning only subscriber host to which the given authentication-policy is applicable.
- Receiving CoA message with the same attributes as currently applicable to the given host will be responded to with an ACK message.
- In case of dual homing (through SRRP), the RADIUS server should send CoA messages to both redundant nodes and this with all corresponding attributes (NAS-port-id with its local meaning to corresponding node).
- In the case of change requests, the node which has the given host active (active-sap for master-sap for SRRP) will process the RADIUS message and reply to RADIUS. The standby node will always reply with a NAK.
- In the case of create requests, the active node (the SAP described by NAS-port-id is "active" or "master"). Both nodes will reply, but the standby will NAK the request.

The properties of an existing RADIUS-authenticated PPPoE session can be changed by sending a Change of Authorization (CoA) message from the RADIUS server. Processing of a CoA is done in the same way as for DHCP hosts, with the exception that only the ESM settings can be changed for a PPPoE session (the Force-Renew attribute is not supported for PPPoE sessions and a Create-Host CoA will always generate a DHCP host.)

For terminating PPPoE sessions from the RADIUS server, the disconnect-request message can be sent from the RADIUS server. This message triggers a shutdown of the PPPoE session. The attributes needed to identify the PPPoE session are the same as for DHCP hosts.

9.2.9.3.1 Change of Authorization using the Tools Command

A CoA can be triggered through the CLI by using a **tools** command that does not require a RADIUS authentication policy. The **tools** command can also be used to spoof a CoA from a configured server for purposes such as testing CoA python scripts. However, when spoofing the CoA from a RADIUS server, the configuration of a RADIUS authentication policy is required.

The **tools** command, **tools**>**perform**>**subscriber-mgmt**>**coa**, supports up to five different VSAs. If more than five VSAs are required, a file with more than five VSAs can be used for execution.

The **tools** command does not support lawful intercept attributes.

SNMP can also trigger the **tools** CoA command. However, SNMP cannot execute the command when it is referencing an on-board flash file. To execute from a file, the file must be non-local, such as using a URL specifying the location of the file on an FTP server.

Only one **tools**>**perform**>**subscriber-mgmt**>**coa** command can be performed at a time. The command must complete execution before processing a new one. If the **tools** command becomes unresponsive, CTRL-c can be used to break out of the CoA. In addition, a failsafe mechanism automatically terminates the **tools** command if it has not completed within a minute.

9.2.9.4 RADIUS-Based Accounting

When a router is configured to perform RADIUS-based accounting, at the creation of a subscriber-host, it will generate an accounting-start packet describing the subscriber-host and send it to the RADIUS accounting server. At the termination of the session, it will generate an accounting-stop packet including accounting statistics for a given host. The router can also be configured to send an interim-accounting message to provide updates for a subscriber-host.

The exact format of accounting messages, their types, and communication between client running on the routers and RADIUS accounting server is described in RFC 2866, *RADIUS Accounting*. The following describes a few specific configurations.

In order to identify a subscriber-host in accounting messages different RADIUS attributes can be included in the accounting-start, interim-accounting, and accounting-stop messages. The inclusion of the individual attributes is controlled by the following commands.

```
configure
    subscr-mgmt
        radius-accounting-policy <name>
            include-radius-attribute
                [no] acct-authentic
                [no] acct-delay-time
                [no] called-station-id
                [no] calling-station-id
                [no] circuit-id
                [no] delegated-ipv6-prefix
                [no] dhcp-vendor-class-id
                [no] framed-interface-id
                [no] framed-ip-addr
                [no] framed-ip-netmask
                [no] framed-ipv6-prefix
                [no] framed-route
                [no] framed-ipv6-route
                [no] ipv6-address
                [no] mac-address
                [no] nas-identifier
```

```
            [no] nas-port
            [no] nas-port-id
            [no] nas-port-type
            [no] nat-port-range
            [no] remote-id
            [no] sla-profile
            [no] sub-profile
            [no] subscriber-id
            [no] tunnel-server-attrs
            [no] user-name
            [no] wifi-rssi
            [no] alc-acct-triggered-reason
            [no] access-loop-options
            [no] all-authorized-session-addresses
            [no] detailed-acct-attributes
            [no] std-acct-attributes
            [no] v6-aggregate-stats
```

RADIUS volume accounting attributes are depending on the type of volume reporting and can be controlled via an **include-radius-attribute** CLI command. Multiple volume reporting types can be enabled simultaneously:

```
config
    subscr-mgmt
        radius-accounting-policy <name>
            include-radius-attribute
                [no] detailed-acct-attributes
                [no] std-acct-attributes
                [no] v6-aggregate-stats
```

where:

detailed-acct-attributes — Report detailed per queue and per policer counters using RADIUS VSAs (enabled by default). Each VSA contains a queue or policer id followed by the stat-mode or 64 bit counter. The VSA's included in the Accounting messages is function of the context (policer or queue, stat-mode, MDA type, and so on):

[26-6527-107] Alc-Acct-I-statmode

[26-6527-127] Alc-Acct-O-statmode

[26-6527-19] Alc-Acct-I-Inprof-Octets-64

[26-6527-20] Alc-Acct-I-Outprof-Octets-64

[26-6527-21] Alc-Acct-O-Inprof-Octets-64

[26-6527-22] Alc-Acct-O-Outprof-Octets-64

[26-6527-23] Alc-Acct-I-Inprof-Pkts-64

[26-6527-24] Alc-Acct-I-Outprof-Pkts-64

[26-6527-25] Alc-Acct-O-Inprof-Pkts-64

[26-6527-26] Alc-Acct-O-Outprof-Pkts-64

[26-6527-39] Alc-Acct-OC-O-Inprof-Octets-64

[26-6527-40] Alc-Acct-OC-O-Outprof-Octets-64

[26-6527-43] Alc-Acct-OC-O-Inprof-Pkts-64

[26-6527-44] Alc-Acct-OC-O-Outprof-Pkts-64

[26-6527-69] Alc-Acct-I-High-Octets-Drop_64

[26-6527-70] Alc-Acct-I-Low-Octets-Drop_64

[26-6527-71] Alc-Acct-I-High-Pack-Drop_64

[26-6527-72] Alc-Acct-I-Low-Pack-Drop_64

[26-6527-73] Alc-Acct-I-High-Octets-Offer_64

[26-6527-74] Alc-Acct-I-Low-Octets-Offer_64

[26-6527-75] Alc-Acct-I-High-Pack-Offer_64

[26-6527-76] Alc-Acct-I-Low-Pack-Offer_64

[26-6527-77] Alc-Acct-I-Unc-Octets-Offer_64

[26-6527-78] Alc-Acct-I-Unc-Pack-Offer_64

[26-6527-81] Alc-Acct-O-Inprof-Pack-Drop_64

[26-6527-82] Alc-Acct-O-Outprof-Pack-Drop_64

[26-6527-83] Alc-Acct-O-Inprof-Octs-Drop_64

[26-6527-84] Alc-Acct-O-Outprof-Octs-Drop_64

[26-6527-91] Alc-Acct-OC-O-Inpr-Pack-Drop_64

[26-6527-92] Alc-Acct-OC-O-Outpr-Pack-Drop_64

[26-6527-93] Alc-Acct-OC-O-Inpr-Octs-Drop_64

[26-6527-94] Alc-Acct-OC-O-Outpr-Octs-Drop_64

[26-6527-108] Alc-Acct-I-Hiprio-Octets_64

[26-6527-109] Alc-Acct-I-Lowprio-Octets_64

[26-6527-110] Alc-Acct-O-Hiprio-Octets_64

[26-6527-111] Alc-Acct-O-Lowprio-Octets_64

[26-6527-112] Alc-Acct-I-Hiprio-Packets_64

[26-6527-113] Alc-Acct-I-Lowprio-Packets_64

[26-6527-114] Alc-Acct-O-Hiprio-Packets_64

[26-6527-115] Alc-Acct-O-Lowprio-Packets_64

[26-6527-116] Alc-Acct-I-All-Octets_64

[26-6527-117] Alc-Acct-O-All-Octets_64

[26-6527-118] Alc-Acct-I-All-Packets_64

[26-6527-119] Alc-Acct-O-All-Packets_64

std-acct-attributes — Report IPv4 and IPv6 aggregated forwarded counters using standard RADIUS attributes (disabled by default):

[42] Acct-Input-Octets

[43] Acct-Output-Octets

[47] Acct-Input-Packets

[48] Acct-Output-Packets

[52] Acct-Input-Gigawords

[53] Acct-Output- Gigawords

v6-aggregate-stats — Report IPv6 aggregated forwarded counters of queues and policers in stat-mode v4-v6 using RADIUS VSAs (disabled by default):

[26-6527-194] Alc-IPv6-Acct-Input-Packets

[26-6527-195] Alc-IPv6-Acct-Input-Octets

[26-6527-196] Alc-IPv6-Acct-Input-GigaWords

[26-6527-197] Alc-IPv6-Acct-Output-Packets

[26-6527-198] Alc-IPv6-Acct-Output-Octets

[26-6527-199] Alc-IPv6-Acct-Output-Gigawords

In addition to accounting-start, interim-accounting, and accounting-stop messages, a RADIUS client on a routers will send also accounting-on and accounting-off messages. An accounting-on message will be sent when a given RADIUS accounting-policy is applied to a given subscriber-profile, or the first server is defined in the context of an already applied policy. The following attributes will be included in such message:

- NAS-identifier
- alc-subscriber-profile-string
- Accounting-session-id
- Event-timestamp

Accounting-off messages will be sent at following events:

- An accounting policy has been removed from a sub-profile.
- The last RADIUS accounting server has been removed from an already applied accounting policy.

These messages contain following attributes:

- NAS-identifier
- alc-subscriber-profile-string
- Accounting-session-id
- Accounting-terminate-cause
- Event-timestamp

In case of dual homing, both nodes will send RADIUS accounting messages for the host, with all attributes as it is locally configured. The RADIUS log files on both boxes need to be parsed to get aggregate accounting data for the given subscriber host regardless the node used for forwarding.

For RADIUS-based accounting, a custom record can be defined to refine the data that is sent to the RADIUS server. Refer to the “Configuring an Accounting Custom Record” in the *7450 ESS, 7750 SR, 7950 XRS, and VSR System Management Guide* for further information.

9.2.9.5 RADIUS Accounting Terminating Cause

The VSA **acct-terminate-cause** attribute provides some termination information. Two additional attributes: [VSA 227] **alc-error-message** and [VSA 226] **alc-error-code** provide additional information in both string and numeric format about the terminating cause of the subscriber session. The full list of error messages and their corresponding error codes can be viewed using the command **tools>dump>aaa>radius-acct-terminate-cause**.

If required, python can alter the content of both VSAs. The following is a python script example where the error codes are remapped from 123 to 8 and from 124 to 17:

```
import alc
import struct

ALU            = 6527
TERM_CAUSE     = 49
ALC_ERROR_CODE = 226

if (alc.radius.attributes.isSet(TERM_CAUSE) and
    alc.radius.attributes.isVSASet(ALU, ALC_ERROR_CODE)):

    error_code = alc.radius.attributes.getVSA(ALU, ALC_ERROR_CODE)
    error_code = struct.unpack('!i', error_code)[0]

    term_cause  = alc.radius.attributes.get(TERM_CAUSE)
    term_cause = struct.unpack('!i', term_cause)[0]

    #print "error code = ", error_code
    #print "term cause = ", term_cause

    # table with mapping from alc-error-code to the standard terminate cause
    # if no mapping will be found, no transformation will be done
    error_map = {
        123 : 8,
        124 : 17
    }

    new_term_cause = error_map.get(error_code, term_cause)
    #print "new term_cause = ", new_term_cause

    alc.radius.attributes.set(TERM_CAUSE, struct.pack('!I', new_term_cause))
```

9.2.9.6 Accounting Modes Of Operation

This section is applicable to the 7750 SR or the 7450 ESS in mixed mode. There are three basic accounting models:

- Per queue-instance
- Per Host

- Per Session

Each of the basic models can optionally be enabled to send interim-updates. Inclusion/exclusion of interim-updates will depend on whether volume based (start/interim-updates/stop) or time based (start/stop) accounting is required.

The difference between the three basic accounting models is in its core related to the processing of the acc-session-id for each model. The differences are related to:

- acct-session-id generation within each model.
- outcome in response to the CoA action relative to the targeted acct-session-id.

The counters for volume-based accounting are collected from queues or policers that are instantiated per sla-profile instance (SPI) on non-HSMDA based hardware or per subscriber on HSMDA based hardware. This is true irrespective of which model of accounting (or combination of models) is deployed. Within accounting context, the SPI on non-HSMDA or subscriber on HSMDA equates to queue-instance.

Table 73 summarizes the key differences between various accounting modes of operation that are supported. Interim-updates for each individual mode can be enabled/disabled via configuration (interim-updates keyword as an extension to the commands that enable three basic modes of accounting). This is denoted by the IU-Config keyword under the 'I-U' column in the table. The table also shows that any two combinations of the three basic models (including their variants for volume/time based accounting) can be enabled simultaneously.

Table 73 **Accounting Modes of Operation**

Accounting Mode	Accounting Entity	START	I-U	STOP	Acct-session-id	Acct-multi-session-id
queue-instance-accounting	queue-instance	X	IU-config	X	X	
	session					
	host					
session-accounting	queue-instance					
	session	X	IU-config	X	X	queue-instance
	host					

Table 73 **Accounting Modes of Operation (Continued)**

Accounting Mode	Accounting Entity	START	I-U	STOP	Acct-session-id	Acct-multi-session-id
host-accounting	queue-instance					
	session					
	host	X	IU-config	X	X	queue-instance
queue-instance-accounting + host-accounting	queue-instance	X	IU-config	X	X	queue-instance
	session					
	host	X	IU-config	X	X	queue-instance
queue-instance-accounting + session-accounting	queue-instance	X	IU-config	X	X	queue-instance
	session	X	IU-config	X	X	queue-instance
	host					
session- accounting + host-accounting	queue-instance					queue-instance
	session	X	IU-config	X	X	
	host	X		X	X	SESSION

Note: Hosts within the targeted CoA entity will be affected as follows:

- If the CoA target is the session, then both constituting members (IPv4 and IPv6) of the dual-stack host will be affected.
- If the CoA target is the queuing-instance, then up to 32 hosts that are sharing that SPI will be affected.

The same principle applies to LI.

The accounting behavior (accounting messages and accounting attributes) in case that the SPI is changed via CoA depends on the accounting mode of operation. On non-HSMDA hardware, the behavior is the following:

- SPI change in conjunction with per queuing instance accounting will trigger a STOP for the old SPI and a START for the new SPI with corresponding counters. Acct-session-id/Acct-Multi-Session-Id will be unique per SPI. Note that Acct-Multi-Session-Id is only generated if per queuing-instance accounting mode of operation is combined with some other mode of operation (host or session).
- SPI change in conjunction with per host or per session accounting (no interim updates for either method) will not trigger any new accounting messages. In other words, SPI change will go unnoticed from the perspective of the accounting server until the host/session is terminated. When the host/session is terminated a STOP will be sent with the VSA carrying the latest SPI name and the acct-multi-session-id attribute of the latest SPI. Acct-session-id will stay the same during the lifetime of the host. Counters are not included in STOP (interim-update not enabled).

  SPI change in conjunction with per host accounting with interim-updates or per session accounting with interim-updates will trigger two interim-update messages:

  – One with the old counters (terminated queues) and the old SPI name VSA. This behavior is similar to the triggered STOP message in per queuing-instance accounting upon SPI change.
  – One with the new counters (new queues instantiated), the VSA carrying the new SPI name and the new acct-multi-session-id referencing the new SPI. This behavior is similar to the triggered START message in per queuing-instance accounting when SPI is changed.

On HSMDA, no START/STOPS are sent since queues are not re-instantiated on ingress or egress.

9.2.9.7 Per Session Accounting

In the per session accounting mode of operation the accounting message stream (START/INTERIM-UPDATE/STOP) is generated per session. The accounting message stream refers to a collection of accounting messages (START/INTERIM-UPDATE/STOP) sharing the same acct-session-id.

- A session is defined for PPPoE hosts for which a state is maintained. The state of the host (single stack or dual-stack) is normally refreshed via PPP keepalives. Each PPPoE host of the same address family (v4 or v6) corresponds to a unique session which is identified by the <session-ID, mac> combination.

In dual-stack PPPoE case, IPv4 and IPv6 hosts are tied to the same (LCP) session. A single authentication request is initiated for such session (triggered by the first host that initiates the session).

For a single stack PPPoE host, the behavior defined in the per session accounting model is indistinguishable from the per host accounting model. The per session accounting model makes difference in behavior only for dual stack PPPoE hosts.

The following are the properties of the Per Session Accounting model:

- A single accounting session ID (acct-session-id) is generated per (PPPoE) session and it can optionally be sent in RADIUS Access-Request message.
- This acct-session-id is synchronized via MCS in dual homing environment.
- The accounting messages (START, INTERIM-UPDATE, STOP) carry the acct-multi-session-id attribute denoting the sla-profile instance with which the session is associated.
- The counters are collected from the queues instantiated through the sla-profile instance. If multiple sessions are sharing the same sla-profile instance, the counters are aggregated. In other words, counters per individual session cannot be extracted from the aggregated count.
- RADIUS triggered changes and LI are applicable per session:
  - Queue/policer RADIUS overrides — Parameters for the referenced queue/policer within the session are changed accordingly.
  - Subscriber aggregate rate limits, scheduler rates and arbiter rates are changed accordingly.
  - CoA DISCONNECT brings down the entire session.
  - LI — Activation based on the session acct-session-id affects the hosts within the session (dual-stack).
  - SLA profile instance change affects all hosts (or sessions) sharing the same sla-profile instance (SPI). If the SPI is changed on a non-HSMDA based MDA, then queues are re-instantiated and counters are reset.
- All applicable IP addresses (v4 and v6 – including all v6 attributes – alc-ipv6-address, framed-ipv6-prefix, delegated-ipv6-prefix) are present in accounting messages for the session.

9.2.9.8 RADIUS Session Accounting with PD as a Managed Route

The Prefix Delegation (PD) prefix is included in the accounting messages using the VSA [99], Framed-IPv6-Route attribute with the string type “pd-host” appended to differentiate it from a regular framed IPv6 route; for example, FRAMED IPV6 ROUTE [99] 39 2001:1000::/64 :: 0 pref 0 type pd-host. PD as a managed route is applicable to both PPP and IPoE sessions and can point either to an IPv4 host or to an IPv6 WAN host.

Table 74 outlines the RADIUS accounting behavior based on the session type and the next-hop host.

Table 74 **Radius Accounting Behavior**

Session Type and Next-Hop Host	RADIUS Accounting Start	RADIUS Accounting Interims	RADIUS Accounting Stop
PPP session with IPv6 PD pointing to IPv4 host as the next hop	A PPP connection will trigger an accounting start	A DHCP NA+PD solicit will trigger an interim update for the PD host with interim reason “delegated-ipv6-prefix-up” and the prefix included in the VSA Framed-IPv6-Route A DHCP PD solicit will trigger an interim update for the PD host with interim reason delegated-ipv6-prefix-up and the prefix included in the VSA framed-ipv6-route **Caveat** • A DHCP PD lease expire will trigger an interim update with interim reason “delegated-ipv6-prefix-down”; however, the VSA framed-ipv6-route is not included	A PPP disconnect with only the IPv4 and IPv6 PD host will trigger an accounting stop with the prefix included in the VSA Framed-IPv6-Route **Caveat** • A PPP disconnect with the IPv4, NA, and PD host without **session-optimized-stop** enabled, will not include the VSA Framed-IPv6-Route

Table 74 **Radius Accounting Behavior (Continued)**

Session Type and Next-Hop Host	RADIUS Accounting Start	RADIUS Accounting Interims	RADIUS Accounting Stop
PPP session with IPv6 PD pointing to IPv6 NA host as the next hop	A PPP connection will trigger an accounting start. It is possible to have a single-stack IPv6-only session	A DHCP NA+PD solicit will trigger an interim update for the PD host with interim reason delegated-ipv6-prefix-up and the prefix included in the VSA framed-ipv6-route **Caveat** • A DHCP PD lease expire will trigger an interim update with interim reason “delegated-ipv6-prefix-down”; however, the VSA FramedIPv6-Route is not included	A PPP subscriber disconnect will trigger an accounting stop with the PD host prefix included in the VSA Framed-IPv6-Route
IPoE session with IPv6 PD pointing to IPv4 host as the next hop	A DHCPv4 or a DHCPv6 request (DHCPv6 always performs NA and PD requests together) will trigger the accounting start	A DHCP PD is always performed together with NA. The PD is not in the start message but is included in the accounting interim update as a part of the host update. If the DHCPv4 lease expires, the interim update will contain the PD prefix in the VSA framed-ipv6-route **Caveat** • A DHCP PD lease expire will trigger an interim update with interim reason “delegated-ipv6-prefix-down”; however, the VSA Framed-IPv6-Route is not included	If only the IPv4 host and PD host remain, the release of the DHCPv4 will trigger an accounting stop with the PD host prefix included in the VSA Framed-IPv6-Route **Caveat** • If the DHCPv4 is released and an IPv6 NA host remains, the IPv6 lease release/expire will be an interim update that will not include the prefix

Table 74 **Radius Accounting Behavior (Continued)**

Session Type and Next-Hop Host	RADIUS Accounting Start	RADIUS Accounting Interims	RADIUS Accounting Stop
IPoE session with IPv6 PD pointing to IPv6 NA host as the next hop	A DHCPv4 or a DHCPv6 request (DHCPv6 always performs NA and PD requests together) will trigger the accounting start. It is possible to have a single-stack IPv6-only session	A DHCP PD is always performed together with NA. The PD is not in the start message but is included in the accounting interim update as a part of the host update. **Caveat** • A DHCP PD lease expire will trigger an interim update with interim reason "delegated-ipv6-prefix-down"; however, the VSA Framed-IPv6-Route is not included	If only the IPv6 subscriber is left, the release of NA will contain the prefix of the PD host **Caveat** • If the DHCPv6 is released and an IPv4 host remains, the IPv6 lease release/expire will be an interim update that will not include the prefix

RADIUS Per Host Accounting

In SR OS, the accounting paradigm is based on SLA profile instances yet this is at odds with traditional RADIUS authentication and accounting which is host-centric. In previous SR OS releases, it was possible to have many hosts sharing a common SLA profile instance, and thus accounting and QoS parameters. Complications would arise with RADIUS accounting because Accounting-Start and Accounting-Stop are a function of sla-profile instance and not the hosts — this meant that some host-specific parameters (like framed-ip-address) would not be consistently included in RADIUS accounting.

Currently, dual-stack subscribers are really two different hosts sharing a single sla-profile instance. A new RADIUS accounting mode has been introduced to support multiple-host environments.

Under accounting-policy, a host-accounting command allows configurable behavior.

9.2.9.9 No Host-Accounting

In prior releases and when no host-accounting is configured, the accounting behavior is as follows:

- A RADIUS accounting start message is sent when the SLA-profile instance is created. It contains accounting (octets/packets) and the framed-ip-address of the host which caused the sla-profile instance to be created.
- Additional hosts may bind to the sla-profile instance at any time, but no additional Accounting messages are sent during these events.
- If the original host disconnects then future Accounting messages will use an IP address of one of the remaining hosts.
- When the final host associated with an sla-profile instance disconnects an Accounting Stop message will be sent.

9.2.9.10 Host-Accounting Enabled

When host-accounting is configured, additional RADIUS accounting messages are created for host activity in addition to messages for common queue accounting. The behavior is as follows:

- A RADIUS accounting start message is sent each time a host is authenticated. It contains the framed-ip-address among other things. It does not contain any octet or packet counts.
- A RADIUS accounting start message is sent each time a sla-profile instance is created.
- Whenever a host disconnects a RADIUS accounting stop message is sent for that host.
- If all host associated with an sla-profile instance disconnect, a RADIUS Accounting Stop message is sent for that instance.

This new behavior means certain AVP may be in either host; sla-profile instance or both accounting records.

Note that interim-acct records are not sent for hosts, only the start- and stop-accting messages. Refer to the *7750 SR and VSR RADIUS Attributes Reference Guide* for attributes that are applicable for RADIUS accounting AVPs.

9.2.9.11 Reduction of Host Updates for Session Accounting Start and Stop

When **host-update** is enabled In session accounting, a dual stack subscriber can generate multiple host update accounting messages at the start and end of a session (for example, one for the IPv4 host and two more for the IPv6 WAN and IPv6 PD hosts). Two features can be used to reduce the number of host update messages per subscriber.

The first feature delays the Start Accounting message by a configurable value and is applicable to both PPPoE and IPoE sessions. The command for configuring this feature is **config>subscr-mgmt>acct-plcy>delay-start-time**. The delay allows the full dual stack address assignment to be completed before triggering the accounting Start message. The Start message reports all the addresses and prefixes assigned to the subscriber at that time. Subsequent new or disconnected hosts will trigger interim host updates if enabled.

The second feature is for PPPoE sessions only and is used to reduce the number of host update messages when a dual stack PPP subscriber disconnects. The command for configuring this feature is **config>subscr-mgmt>sub-prof>rad-acct>session-optimized-stop**. A single accounting Stop message containing all the addresses and prefixes for the subscriber at the time is generated.

9.2.9.12 Accounting Interim Update Message Interval

The interval between two RADIUS Accounting Interim Update messages can be configured in the RADIUS accounting policy with the **update-interval** command, for example:

```
config
 subscr-mgmt
        radius-accounting-policy "acct-policy-1" create
            update-interval 60
            update-interval-jitter absolute 600
```

A RADIUS specified interim interval (attribute [85] Acct-Interim-Interval) overrides the CLI configured value.

By default, a random delay of 10% of the configured **update-interval** is added to the update-interval between two Accounting Interim Update messages. This jitter value can be configured with the **update-interval-jitter** to an absolute value in seconds between zero and 3600. The effective maximum random delay value is the minimum value of the configured absolute jitter value and 10% of the configured **update-interval**.

A value of zero will send the Accounting Interim Update message without introducing an additional random delay.

9.2.9.13 CoA Triggered Accounting Interim Update

The vendor-specific attribute (VSA) [228], Alc-Triggered-Acct-Interim, can be used in a Change of Authorization message to trigger an interim accounting message. This feature requires the accounting mode to have interim updates enabled. You can enable interim updates using, the **config>subscr-mgmt>radius-acct-plcy>host-accounting interim-update** command. The VSA can hold a string of up to 247 characters. The accounting interim echoes this string in the interim message under the same Alc-Triggered-Acct-Interim VSA along with Alc-Acct-Triggered-Reason = CoA-triggered. If the VSA is left blank, it will still trigger the accounting interim message with Alc-Acct-Triggered-Reason = CoA-triggered (18), but without the Alc-Triggered-Acct-Interim attribute. If the subscriber session has multiple accounting policies or modes enabled, multiple interim messages will be generated.

Some CoAs, such as SLA profile or sub-profile changes, will trigger accounting update messages to be generated automatically. These CoAs can automatically generate one or more accounting interim messages. If these CoAs also include the Alc-Triggered-Acct-Interim VSA, no additional interim accounting messages will be generated. The last automatically generated accounting interim message will contain two reasons:

- the reason for the triggered interim message (such as an SLA start)
- the CoA-triggered (18) Alc-Triggered-Acct-Interim attribute that will be echoed on the triggered accounting interim message if the VSA is not empty

9.2.9.14 Class Attribute

The RADIUS class attribute helps to aid in user identification.

User identification is used to correlate RADIUS accounting messages with the given user. During the authentication process, the RADIUS authentication server inserts a class attribute into the RADIUS authenticate response message and the router echoes this class attribute in all RADIUS accounting messages.

The 7750 SR can store up to six class attributes for both RADIUS and NASREQ. Each class VSA or AVP can have a maximum of 253 characters. If the VSA or AVP contains more than 253 characters, only the first 253 characters will be stored. If there are more than six VSAs or AVPs, only the first six will be stored. This functionality is also applicable to RADIUS authentication via the ISA.

9.2.9.15 User Name

The user-name, which is used for user authentication (user-name attribute in RADIUS authentication request), can be included in RADIUS accounting messages. Per RFC 2865, when a RADIUS server returns a (different) user-name attribute, the changed user name will be used in accounting and not the originally sent user name.

9.2.9.16 Accounting-On and Accounting Off

For RADIUS servers configured in a RADIUS server policy, the accounting on/off behavior is controlled via the **acct-on-off** command in the radius-server-policy.

By default, no Accounting-On or Accounting-Off messages are sent (**no acct-on-off**).

With the **acct-on-off** command configured in the radius-server-policy:

- An Accounting-On is sent for the following:
  - When the system is powered on.
  - After a system reboots.
  - When the **acct-on-off** command is added to the **radius-server-policy** configuration.
  - User triggered via CLI: tools perform aaa acct-on
- An Accounting-Off is sent for the following:
  - Before a user initiated system reboot.
  - When the **acct-on-off** command is removed from the **radius-server-policy** configuration.
  - User triggered via CLI: tools perform aaa acct-off.

The Accounting-On or Accounting-Off message is sent to the servers configured in the radius-server-policy, following the configured access-algorithm until an Accounting Response is received. If the first server responds, no message is sent to the other servers.

The Accounting-On message is repeated until an Accounting Response message is received from a RADIUS server: If after the configured retry/timeout timers for each RADIUS server in the radius-server-policy no response is received then the process starts again after a fixed one minute wait interval.

The Accounting-Off message is attempted once: If after the configured retry/timeout timers for each RADIUS server in the radius-server-policy no response is received then no new attempt is made.

It is possible to block a radius-server-policy until an Accounting Response is received from one of the RADIUS servers in the radius-server-policy that acknowledges the reception of an Accounting-On. The radius-server-policy cannot be used by applications for sending RADIUS messages until the state becomes “Not Blocked”. This is achieved with the optional “oper-state-change” flag, for example:

```
config
    aaa
        radius-server-policy "aaa-server-policy-1" create
            acct-on-off oper-state-change
            servers
                router "Base"
                server 1 name "server-1"
            exit
        exit
    exit
```

If multiple radius-server-policies are in use for different applications (for example, authentication and accounting) and an Accounting-On must be send for only one radius-server-policy, it is possible to tie the acct-on-off states of both policies together using an acct-on-off-group. With this configuration, it is possible to block the authentication servers until the accounting servers are available. An acct-on-off-group can be referenced by:

- a single radius-server-policy as controller: the acct-on-off oper-state of the acct-on-off-group is set to the acct-on-off oper-state of the radius-server-policy (acts as master)
- multiple radius-server-policies as monitor: the acct-on-off oper-state of the radius-server-policy is inherited from the acct-on-off oper-state of the acct-on-off group. (acts as a slave)

```
config
    aaa
        acct-on-off-group "group-1" create
            description "Grouping of radius-server-policies acct-on-off"
        exit
```

```
radius-server-policy "aaa-server-policy-1" create
    acct-on-off oper-state-change group "group-1"
    servers
        router "Base"
        server 1 name "server-1"
    exit
exit
radius-server-policy "aaa-server-policy-2" create
    acct-on-off monitor-group "group-1"
    servers
        router "Base"
        server 1 name "server-2"
    exit
exit
```

It is possible to force an Accounting-On or Accounting-Off message for a radius-server-policy with acct-on-off enabled using following CLI commands:

tools perform aaa acct-on [**radius-server-policy** *policy-name*] [**force**]

tools perform aaa acct-off [**radius-server-policy** *policy-name*] [**force**] [**acct-terminate-cause** *number*]

If an Accounting-On was sent to the radius-server-policy and it was acknowledged with an Accounting Response then a new Accounting-On can only be sent with the “force” flag.

If an Accounting-Off was sent to the radius-server-policy and it was acknowledged with an Accounting Response then a new Accounting-Off can only be sent with the “force” flag. The Acct-Terminate-Cause value in the Accounting-Off can be overwritten.

Use the following CLI command to display the Accounting On/Off information for a radius-server-policy:

```
# show aaa radius-server-policy "aaa-server-policy-3" acct-on-off
===============================================================================
RADIUS server policy "aaa-server-policy-3" AcctOnOff info
===============================================================================
Oper state                  : on
Session Id                  : 242FFF0000008F512A3985
Last state change           : 02/24/2013 16:06:41
Trigger                     : startUp
Server                      : "server-1"
===============================================================================
```

The operational state provides following state information: The sending of the Accounting-On or Accounting-Off message is ongoing (sendAcctOn, SendAcctOff), is successfully responded (on, off) or no response received (OffNoResp).

The Session-Id is a unique identifier for each RADIUS server policy accounting Accounting-On/Accounting-Off sequence.

The Trigger field shows what triggered the Accounting On or Accounting Off message. If the radius-server-policy is part of an acct-on-off group then the group name is shown in brackets.

The Server field shows which server in the RADIUS server policy responded to the Accounting-On or Accounting-Off message.

To display the acct-on-off state of a radius-server-policy, use the command, for example:

```
# show aaa radius-server-policy "aaa-server-policy-3"
===============================================================================
RADIUS server policy "aaa-server-policy-3"
===============================================================================
Description                 : (Not Specified)
Acct Request script policy  : script-policy-1
Auth Request script policy  : script-policy-1
Accept script policy        : script-policy-1
Acct-On-Off                 : Enabled (state Blocked)
-------------------------------------------------------------------------------
RADIUS server settings
-------------------------------------------------------------------------------
Router                      : "Base"
Source address              : (Not Specified)
Access algorithm            : direct
Retry                       : 3
Timeout (s)                 : 5
Hold down time (s)          : 30
Last management change      : 02/20/2013 13:32:05
===============================================================================
===============================================================================
Servers for "aaa-server-policy-3"
===============================================================================
Idx Name                              Address         Port        Oper State
                                                      Auth/Acct
-------------------------------------------------------------------------------
1   server-3                          172.16.1.10     1812/1813   unknown
===============================================================================
```

The Acct-On-Off field indicates if the sending of Accounting-On and Accounting-Off messages is enabled or disabled. If enabled, the oper-state is displayed: state Blocked or state Not Blocked. When Blocked, the radius-server-policy cannot be used to send RADIUS messages.

To display acct-on-off-group information, use following command, for example:

```
# show aaa acct-on-off-group "group-1"
===============================================================================
Acct-On-Off-Group Information
===============================================================================
acct on off group name              : group-1
```

```
  - controlling Radius-Server-policy :
        aaa-server-policy-1
  - monitored by Radius-Serer-policy :
        aaa-server-policy-2

-------------------------------------------------------------------------------
Nbr of Acct-on-off-groups displayed : 1
-------------------------------------------------------------------------------
===============================================================================
```

9.2.9.17 RADIUS Accounting Message Buffering

When all servers in a radius-server-policy are unreachable, it is possible to buffer the Accounting Stop and Accounting Interim-Update messages for up to 25 hours. When a RADIUS server becomes reachable again then the messages in the buffer are retransmitted.

RADIUS Accounting message buffering parameters can be configured per message type, for example:

```
config
    aaa
        radius-server-policy "aaa-server-policy-1" create
            servers
                router "Base"
                buffering
                    acct-interim min 60 max 3600 lifetime 12
                    acct-stop min 60 max 3600 lifetime 12
                exit
                server 1 name "server-1"
            exit
        exit
    exit
```

When RADIUS accounting message buffering is enabled:

1. The message is stored in the buffer, a lifetime timer is started and the message is sent to the RADIUS server.
2. If after retry*timeout seconds no RADIUS accounting response is received for the Accounting Interim Update or Accounting Stop then a new attempt to send the message is started after minimum [(min-val*2n), max-val] seconds.
3. Repeat step 2 until:
    a. RADIUS accounting response is received, or
    b. the lifetime of the buffered message expires, or
    c. (if the buffered message is an Accounting Interim-Update only) A new Accounting Interim-Update or an Accounting Stop or for the same accounting session-id and radius-server-policy is stored in the buffer, or

d. the message is manually purged from the message buffer with a clear command

- The message is purged from the buffer as shown in Figure 69.

Figure 69 **Purging Message from Buffer**

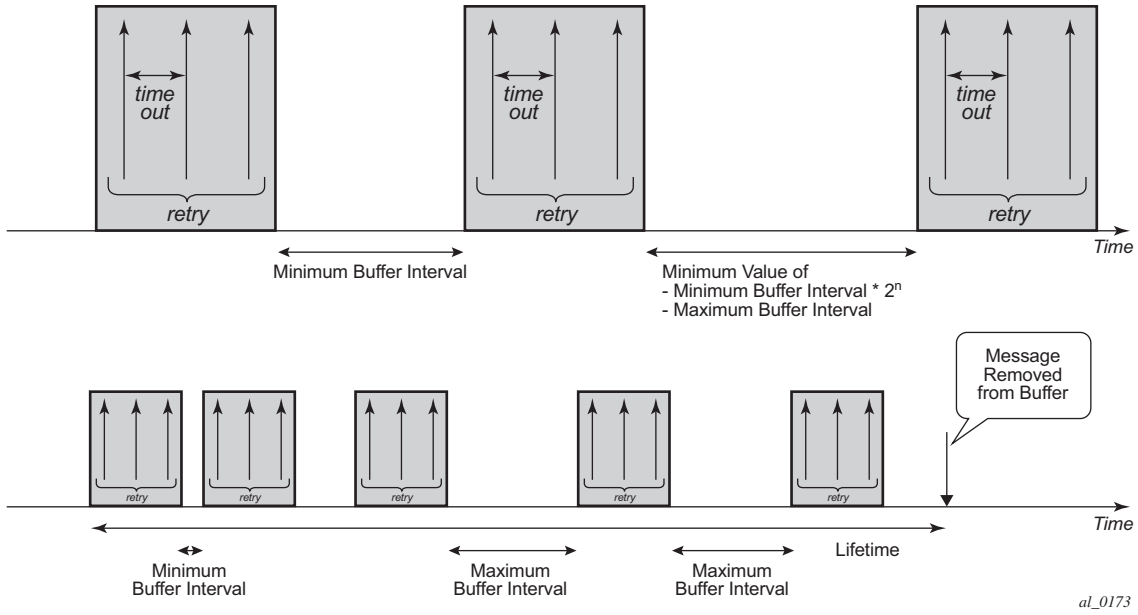

When Accounting Interim-Update message buffering is enabled, it is recommended to also enable Accounting Stop message buffering. This will guarantee the message ordering per accounting session.

Use following clear command to manually delete messages from the RADIUS accounting message buffer:

 # clear aaa radius-server-policy *policy-name* **msg-buffer** [**acct-session-id** *acct-session-id*]

When specifying the Acct-Session-Id, only that specific message will be deleted from the message buffer. If no Acct-Session-Id is specified, all messages for that radius-server-policy are deleted from the message buffer.

Use the following **show** commands to display the RADIUS accounting message buffer statistics:

```
# show aaa radius-server-policy "aaa-server-policy-1" msg-buffer-stats
===============================================================================
RADIUS server policy "aaa-server-policy-1" message buffering stats
===============================================================================
buffering acct-interim      : enabled
  min interval (s)          : 60
```

```
  max interval (s)          : 3600
  lifetime (hrs)            : 12
buffering acct-stop         : enabled
  min interval (s)          : 60
  max interval (s)          : 3600
  lifetime (hrs)            : 12

Statistics
-------------------------------------------------------------------------------
Total acct-stop messages in buffer                      : 0
Total acct-interim messages in buffer                   : 5
Total acct-stop messages dropped (lifetime expired)     : 0
Total acct-interim messages dropped (lifetime expired)  : 0
Last buffer clear time                                  : N/A
Last buffer statistics clear time                       : N/A
-------------------------------------------------------------------------------
===============================================================================
```

Use following **clear** command to reset the RADIUS accounting message buffer statistics:

 # clear aaa radius-server-policy *policy-name* **statistics msg-buffer-only**

Use following **tools** commands to display the RADIUS accounting message buffer content:

 # tools dump aaa radius-server-policy *policy-name* **msg-buffer [session-id** *acct-session-id***]**

For example:

```
# tools dump aaa radius-server-policy "aaa-server-policy-1" msg-buffer
===============================================================================
RADIUS server policy "aaa-server-policy-1" message buffering
===============================================================================
message type Acct-Session-Id                              remaining lifetime
-------------------------------------------------------------------------------
acct-interim 242FFF0000009A512B36FC                       0d 11:58:54
acct-interim 242FFF0000009B512B36FC                       0d 11:58:48
acct-interim 242FFF0000009C512B36FC                       0d 11:58:30
acct-interim 242FFF0000009D512B36FC                       0d 11:58:29
acct-interim 242FFF0000009E512B36FC                       0d 11:59:05
-------------------------------------------------------------------------------
No. of messages in buffer: 5
===============================================================================
```

When specifying the Acct-Session-Id, the message details are displayed.

9.2.9.18 Multiple Accounting Policies

The subscriber profile allows the user to configure a primary accounting policy with an additional accounting policy. The accounting policies are independent of each other and each policy has its own accounting mode, update interval, and include attributes. The RADIUS VSA [85] Acct-Interim-Interval attribute changes both the primary and the duplicate accounting interim update interval.

9.2.9.19 Sending an Accounting Stop Message upon a RADIUS Authentication Failure of a PPPoE Session

In scenarios where RADIUS authentication is used for PPPoE sessions, an accounting stop message can be generated to notify the RADIUS servers in case of an authentication failure.

The failure events are categorized in three categories:

- "**on-request-failure**" — All failure conditions between the sending of an Access-Request and the reception of an Access-Accept or Access-Reject.
- "**on-reject**" — When an Access-Reject is received.
- "**on-accept-failure**" — All failure conditions that appear after receiving an Access-Accept and before successful instantiation of the host or session.

Each of the categories can be enabled separately in the RADIUS authentication policy.

In the Enhanced Subscriber Management (ESM) model, the RADIUS accounting server is found after authentication and host identification as part of the subscriber profile configuration. To report authentication failures to accounting servers, an alternative RADIUS accounting policy configuration is required: local user database pre-authentication can provide the RADIUS authentication policy to be used for authentication and the RADIUS accounting policy to be used for authentication failure reporting. A duplicate RADIUS accounting policy can be specified if the accounting stop resulting from a RADIUS authentication failure must also be sent to a second RADIUS destination.

```
configure
    subscriber-mgmt
        local-user-db "ludb-1" create
            ppp
                match-list username
                host "default" create
                    auth-policy "auth-policy-1"
                    password ignore
```

```
                acct-policy "acct-policy-1" duplicate "acct-policy-2"
                no shutdown
            exit
        exit
        no shutdown
    exit
    authentication-policy "auth-policy-1" create
        pppoe-access-method pap-chap
        include-radius-attribute
           - - - snip - - -
        exit
        send-acct-stop-on-fail on-request-failure on-reject on-accept-failure
        radius-server-policy "aaa-server-policy-1"
    exit
    radius-accounting-policy "acct-policy-1" create
        - - - snip - - -
        radius-server-policy "aaa-server-policy-1"
    exit
    radius-accounting-policy "acct-policy-2" create
        - - - snip - - -
        radius-server-policy "aaa-server-policy-2"
    exit
```

To enable local user database pre-authentication, use the **user-db** configuration in the capture SAP and in the group interface. For example:

```
configure
    service
        vpls 10 customer 1 create
            sap 1/1/1:1.* capture-sap create
                trigger-packet pppoe
                pppoe-policy "ppp-policy-1"
                pppoe-user-db "ludb-1"
            exit
            no shutdown
        exit
        ies 1000 customer 1 create
            subscriber-interface "sub-int-1" create
               - - - snip - - -
                group-interface "group-int-1-1" create
                    - - - snip - - -
                    pppoe
                        policy "ppp-policy-1"
                        user-db "ludb-1"
                        no shutdown
                    exit
                exit
            exit
            no shutdown
        exit
```

9.3 Enhanced Subscriber Management Overview

9.3.1 Enhanced Subscriber Management Basics

In residential broadband networks numerous subscribers can be provisioned that can require significant changes on a daily basis. Manually configuring the applicable parameters for each subscriber would be prohibitive. The Nokia 7450 ESS and 7750 SR have been designed to support fully dynamic provisioning of access, QoS and security aspects for residential subscribers using DHCP to obtain an IP address. Enabling Enhanced Subscriber Management drastically reduces the configuration burden.

Enhanced Subscriber Management in the 7450 ESS and 7750 SR supports many vendor's access nodes and network aggregation models, including VLAN per customer, per service or per access node.

9.3.1.1 Standard and Enhanced Subscriber Management

The system can switch between standard and enhanced subscriber management modes on a per SAP basis. The Enhanced Subscriber Management mode is supported on the SR-7 and SR-12 chassis and on the ESS-7 chassis.

Some functions are common between the standard and enhanced modes. These include DHCP lease management, static subscriber host definitions and anti-spoofing. While the functions of these features may be similar between the two modes, the behavior is considerably different.

- Standard mode — The system performs SLA enforcement functions on a per SAP basis, that is, the attachment to a SAP with DHCP lease management capabilities. The node can authenticate a subscriber session with RADIUS based on the MAC address, the circuit-id (from Option 82) or both. It will then maintain the lease state in a persistent manner. It can install anti-spoofing filters and ARP entries based on the DHCP lease state. Static subscriber hosts are not required to have any SLA or subscriber profile associations and are not required to have a subscriber identification string defined.
- Enhanced mode — When enabled on a SAP, the system expands the information it stores per subscriber host, allowing SLA enforcement and accounting features on a per subscriber basis. The operator can create a subscriber identification policy that will include a URL to a user-space script that assists with the subscriber host identification process.

- A subscriber host is identified by a subscriber identification string instead of the limited Option 82 values (although, the identification string is normally derived from string manipulation of the Option 82 fields). A subscriber identification policy is used to process the dynamic host DHCP events to manage the lease state information stored per subscriber host. The static subscriber hosts also must have subscriber identification strings associations to allow static and dynamic hosts to be grouped into subscriber contexts.
- Further processing by the subscriber identification policy derives the appropriate subscriber and SLA profiles used to define the hierarchical virtual schedulers for each subscriber and the unique queuing and filtering required for the hosts associated with each subscriber
- The SLA profile information is used to identify which QoS policies and which queues/policers, and also which egress hierarchical virtual schedulers, will be used for each subscriber host (dynamic or static).
- The system performs SLA enforcement functions on a per subscriber SLA profile instance basis. SLA enforcement functions include QoS (classification, filtering and queuing), security (filtering), and accounting.

When the enhanced mode is enabled on a SAP (see Subscriber SAPs), first, the router ensures that existing configurations on the SAP do not prevent proper enhanced mode operation. If any one of the following requirements is not met, enhanced mode operation is not allowed on the SAP:

- Anti-spoofing filters must be enabled and configured as IP+MAC matching.
- Any existing static subscriber hosts must have:
  - An assigned subscriber identification string.
  - An assigned subscriber profile name.
  - An assigned SLA profile name.
- The system must have sufficient resources to create the required SLA profile instances and schedulers.

When the router successfully enables the enhanced mode, the current dynamic subscriber hosts are not touched until a DHCP message event occurs that allows re-population of the dynamic host information. Thus, over time, the dynamic subscriber host entries are moved from SAP-based queuing and SAP-based filtering to subscriber-based queuing and filtering. In the event that a dynamic host event cannot be processed due to insufficient resources, the DHCP ACK message is discarded and the previous host lease information is retained in the system.

9.3.1.1.1 Subscriber Management Definitions

Subscriber

A subscriber is typically defined by a unique subscriber identifier to which an assortment of polices (or subscriber profile) can be applied. A subscriber typically (but not always) maps into a VLAN, a VPI/VCI pair, an “ifentry” (a logical interface such as a SAP), a (source) MAC or IP address or a physical port, which uniquely identify a billable entity for the service provider.

Subscriber Management

The management of all services, policies, AAA functions and configurations that relate to the concept of a subscriber. Subscriber management can be configured in a variety of ways, but it is critical that subscriber management integrates seamlessly with element and service management across the broadband infrastructure, via for instance, the 5750 Subscriber Services Controller (SSC). Subscriber management can also be implemented through CLI or scripted commands at the platform level, whereby a network administrator would manually configure the set of QoS, security, AAA or anti-spoofing functions that relate to a particular billable entity or subscriber. Subscriber management is typically centralized and highly integrated with the element, services and middleware management functions for streamlined management, flowthrough provisioning, and accelerated service activation, with minimized operating expenditures.

Subscriber Policy Enforcement

Is the set of actual enforcement functions that are implemented relative to a given subscriber, possibly at multiple enforcement points in the infrastructure and as a result of a match between the subscriber profile which was defined by the subscriber management suite (5750 SSC) and actual traffic patterns. Examples include for instance, the shaping, policing or rate limiting of traffic or the traffic of a given subscriber being dropped because it matched or violated any specific rule (packet with a mismatch between MAC and IP address suggesting an address spoof for instance)

Subscriber SAPs

A subscriber SAP is a service access point (SAP) where enhanced subscriber management is active. Enhanced subscriber management must be explicitly enabled on a per-SAP basis with the CLI **sub-sla-mgmt** command.

A subscriber SAP can be used by a single subscriber or support multiple subscribers simultaneously. Each subscriber can be represented by one or multiple subscriber hosts on the subscriber SAP. If enhanced subscriber management is enabled on a SAP, any configured QoS and IP filter policies defined on the SAP are ignored. A subscriber SAP must refer to an existing subscriber identification policy.

Hosts and Subscribers

A host is a device identified by a unique combination of IP address and MAC address. Typically, the term “subscriber host” is used instead of the “host”.

A host can be an end-user device, such as a PC, VoIP phone or a set top box, or it can be the user’s Residential Gateway (RGW) if the RGW is using Network Address Translation (NAT).

Each subscriber host must be either statically provisioned or dynamically learned by the system. The host’s IP address + MAC address are populated in the subscriber host table on the appropriate SAP to allow packets matching the IP address and MAC address access to the provider’s network.

- A dynamic subscriber host is dynamically learned by the system through the DHCP snooping or relay process. Each subscriber SAP created on the system is configured (using the lease-populate command) to monitor DHCP activity between DHCP clients reached through the SAP and DHCP servers. DHCP ACKs from the DHCP server are used to determine that a certain IP address is in use by a specific DHCP client. This client IP address association is treated by the system as a dynamic subscriber host.
- When it is not possible to dynamically learn a subscriber host through DHCP, a static subscriber host can be created directly on a subscriber SAP. Since a subscriber identification policy is not applicable to static subscriber hosts, the subscriber identification string, subscriber profile and SLA profile must be explicitly defined with the hosts IP address and MAC address.

A subscriber (in the context of the router) is a collection of hosts getting common (overall) treatment. It is expected that this group of hosts originate from the same site and all hosts of a subscriber are reached by the same physical path (such as a DSL port).

Once a subscriber host is known by the system, it is associated with a subscriber identifier and an SLA profile instance. Subscriber hosts with a common subscriber identifier are considered to be owned by the same subscriber.

Depending on the network model, hosts associated with a single subscriber can be associated with a single subscriber SAP or spread across multiple subscriber SAPs on the same port.

9.3.1.1.2 Subscriber Identification Policy

The subscriber identification policy contains the URL definitions for the Programmable Subscriber Configuration Policy (PSCP) scripts used for DHCP ACK message processing. Up to three URLs can be defined per subscriber identification policy. These are designated as primary, secondary and tertiary. Each URL can be individually enabled or disabled. Only one script (the URL with the highest priority active script) is used at any one time to process DHCP ACK messages. If the system detects an error with a specified script, the URL is placed in an operationally down state. If the script is shutdown, it is placed in an administratively down state. A script that is operationally or administratively down is considered inactive. The system automatically reverts to the highest priority active script. If a script becomes operationally down, it must be cycled through the administratively down then administratively up states for the system to attempt to reactivate the script.

Multiple subscriber identification policies are provided for the event that access nodes (such as DSLAMs) from different vendors are attached to the same router. Each policy's active script can be explicitly defined to process the various DHCP message formats or idiosyncrasies of each vendor.

If a script is changed, it must be reloaded by disabling and re-enabling any URL which refers to the changed script (a **shutdown** command followed by a **no shutdown** command).

Each subscriber identification policy can also contain a subscriber profile map and/or an SLA profile map. The subscriber profile map creates a mapping between the sub-profile-strings returned from the active script with an existing subscriber profile name. The SLA profile map is used to create a mapping between the sla-profile-strings returned from the active script with an existing SLA profile name.

The subscriber identification policy is designed to accept a DHCP ACK message destined for a subscriber host and return up to three string values to the system;

- The subscriber identification string (mandatory)
- The subscriber profile string (optional)
- The SLA profile string (optional).

These strings are used to derive the subscriber profile and the SLA profile to be used for this host See Using Scripts for Dynamic Recognition of Subscribers.

9.3.1.1.3 Subscriber Identification String

Subscribers are managed by the router through the use of subscriber identification strings. A subscriber identification string uniquely identifies a subscriber.

The subscriber identification string is the index key to any entry in the active subscriber table, and thus must always be available. It is derived as follows:

- For dynamic hosts, the subscriber identification string is derived from the DHCP ACK message sent to the subscriber host.
  - The DHCP ACK message is processed by a subscriber identification script which has the capability to parse the message into an alternative ASCII string value.
  - If enhanced subscriber management is disabled, the default value for the string is the content of the Option 82 circuit-id and remote-id fields interpreted as an octet string.
- For static hosts, the subscriber identification string must be explicitly defined with each static subscriber host.

When multiple hosts are associated with the same subscriber identification string, they are considered to be host members of the same subscriber. Hosts from multiple SAPs can be members of the same subscriber, but for proper virtual scheduling to be performed all hosts of a subscriber must be active on the same IOM.

When the first host (either dynamic or static) is created with a certain subscriber identification string, an entry is created in the active subscriber table. The entries are grouped by their subscriber identification string.

9.3.1.1.4 Subscriber Profile

The subscriber profile is a template which contains those hierarchical QoS (HQoS) and accounting settings which are applicable to all hosts belonging to the same subscriber. These include:

- Ingress and egress scheduler policy HQoS
- Accounting policy
- RADIUS accounting policy

Subscribers are either explicitly mapped to a subscriber profile template or are dynamically associated with a subscriber profile.

Attempting to delete any subscriber profile (including the profile named ‘default’) while in use by an existing active subscriber will fail.

9.3.1.1.5 SLA Profile

For the purpose of supporting multiple service types (such as high speed Internet (HSI), voice over IP (VoIP), video on demand (VoD) and Broadcast TV) for a single subscriber, the hosts associated with a subscriber can be subdivided into multiple SLA profiles.

The SLA profile contains those QoS and security settings which are applicable to individual hosts. An SLA profile acts like a template and can be used by many subscribers at one time. Settings in the SLA profile include:

- Egress and ingress QoS settings
- Egress scheduler policy HQoS
- Egress and ingress IP filters
- Host limit

If the SLA profile does not explicitly define an ingress or egress QoS policy, the default SAP ingress or default SAP egress QoS policy is used.

See Determining the SLA Profile for information on how the SLA profile is determined for dynamic hosts.

9.3.1.1.6 Explicit Subscriber Profile Mapping

An explicit mapping of a subscriber identification string to a specific subscriber profile can be configured.

An explicit mapping overrides all default subscriber profile definitions while processing a DHCP ACK. In an environment where dynamic and static hosts coexist in the context of a single subscriber, care will be taken to not define a subscriber profile in the explicit subscriber map that conflicts with the subscriber profile provisioned for the static host(s). If such a conflict occurs, the DHCP ACKs will be dropped.

An explicit mapping of a subscriber identification string to the subscriber profile name ‘default’ is not allowed. However, it is possible for the subscriber identification string to be entered in the mapping table without a defined subscriber profile which can result in the explicitly defined subscriber to be associated with the subscriber profile named ‘default’.

Attempting to delete a subscriber profile that is currently defined in an explicit subscriber identification string mapping will fail.

The explicit mapping entries can be removed at any time.

9.3.2 ESM for IPv6

ESM for IPv6 is supported on 7750 SR chassis or in 7450 ESS chassis operating in mixed mode. ESM for IPv6 is supported with RADIUS as the backend authentication and authorization mechanism.

9.3.2.1 Models

9.3.2.1.1 PPPoE Host

For PPPoE, the ESR suggests the IPv6CP protocol to the client during the session setup phase if the appropriate attributes have been returned by the RADIUS server on authentication. The RADIUS attribute that indicates the setup of a PPPoE host is Framed-IPv6-Prefix, which should contain a /64 prefix for the client.

When a PPPoE host has successfully completed the IPv6CP negotiation, the ESR will transmit a Router Advertisement to the PPPoE host containing the suggested prefix and any other options that are configured. The client may use this information to pick one or more addresses from the suggested prefix; all addresses within the prefix are forwarded towards the client.

Alternatively, the Recursive DNS Server (RDNSS) Option as defined in RFC 6106, I*Pv6 Router Advertisement Options for DNS Configuration*, can be included in IPv6 Router Advertisements for DNS name resolution of IPv6 SLAAC hosts. The following CLI command includes the DNS info in IPv6 Router Advertisements for SLAAC hosts and sets the RDNSS lifetime:

```
config>service>ies>sub-if>grp-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv

[no] dns-options

   [no] include-dns    - Set/reset inclusion of the RDNSS server
                         option 25 on this group-interface
   [no] rdnss-lifetime – Maximum time the RDNSS address is valid
                         in this group-interface
```

The source for DNS information to be included in Router Advertisements for IPv6 SLAAC hosts, can be (listed in priority order):

1. Local User Database IPv6 options:

   configure subscriber-mgmt local-user-db <ludb-name> dhcp|ppp host <host-name> options6 dns-server <ip-address> [<ip-address>...(up to 4 max)]

2. RADIUS attributes [26-6527-105] Alc-Ipv6-Primary-Dns and [26-6527-106] Alc-Ipv6-Secondary-Dns
3. Default IPv6 DNS Server configured at the group interface:

   configure service ies|vprn <svc-id> subscriber-interface <sub-int-name> ipv6 default-dns <ipv6-address> [secondary <secondary-ipv6-address>]

Note: A default IPv6 server configuration at the group interface is a last resort IPv6 DNS info that can be used for IPoEv6 hosts (IA_NA, IA_PD and SLAAC) and PPPoEv6 hosts (IA_NA, IA_PD and SLAAC).

9.3.2.1.2 PPPoE RG

Initially, a PPPoE RG follows the same procedure as a PPPoE host: the ESR receives a prefix from RADIUS (in this case through a Delegated-IPv6-Prefix attribute), which is used as a trigger to suggest the IPv6CP protocol to the client. The prefix that is suggested to the client should have the same prefix length as configured under the subscriber interface ipv6 node (delegated-prefix-length). This length should be between 48 and 64 bits, inclusive.

After the IPv6CP protocol has completed, however, the client should run the DHCPv6 protocol over its PPPoE tunnel to receive a Delegated Prefix (IA_PD) and optionally IPv6 DNS server information. This Delegated Prefix can then be subdivided by the client and distributed over its downstream interfaces. During DHCPv6, no extra RADIUS request will be made; the information is stored during the initial (PPPoE or PPP) authentication until the client starts DHCPv6.

Only after DHCPv6 has completed, the IPv6 subscriber host will be instantiated and the ESR will start sending Router Advertisements (if configured.) The router advertisements will not contain any prefix information, which has already been provided by DHCPv6, but it is used as an indication to the client that its default gateway should be the ESR.

9.3.2.1.3 IPoE Host/RG

Similar to an IPv4 DHCP client, a DHCPv6 client is authenticated at its Solicit message, where it can request one or more addresses or prefixes. The address and prefix types supported are IA_NA (Non-Temporary Address) through the Alc-IPv6-Address RADIUS attribute and IA_PD (Delegated Prefix) through the Delegated-IPv6-Prefix attribute. Contrary to the IPv4 case, the ESR will always reply to a DHCPv6 request because the client may request more than one address or prefix simultaneously and not all of the requests may be honored.

The DHCPv6 protocol handling and Router Advertisement behavior are similar to the PPPoE RG case above, with the exception that for an IA_NA address, the entire /64 prefix containing the address is allocated to the client.

For SLAAC prefix assignment, authentication is triggered on router-solicit message. The SLAAC prefix can be assigned statically or dynamically. For a static SLAAC prefix, frame-ipv6-prefix, RADIUS attribute is used. For dynamic SLAAC prefix assignment from a local pool, Alc-slaac-ipv6-pool, RADIUS attribute is used.

9.3.2.2 Setup

IPv6 ESM hosts are only supported in the Routed CO model (both VPRN and IES).

At the IPv6 node under the subscriber interface level, the length of the prefixes that are offered is defined through the delegated-prefix-length option. This setting is fixed for the subscriber interface and can not be changed once subscriber prefixes are defined.

Subscriber prefixes define the ranges of addresses that are offered on this subscriber interface. By default, only these subscriber prefixes are exported to the routing protocols to keep the routing tables small. There are three types of subscriber interfaces:

- wan-host — A range of prefixes that are assigned to PPPoE hosts and as DHCPv6 IA_NA addresses. These prefixes are always /64.
- pd — A range of prefixes that are assigned as DHCPv6 IA_PD prefixes for DHCPv6 IPoE clients and for PPPoE RGs. The length of these prefixes is defined by the delegated-prefix-length.
- both — When both 'wan-host' and 'pd' are defined, the subscriber prefix is a range that can be used for both previous types. However, the delegated-prefix-length is restricted to /64 in this case.

The subscriber interface prefix can also be provisioned via RADIUS. The RADIUS VSA Alc-IPv6-Sub-If-Prefix requires a prefix and the prefix type. The prefix type can be **pd**, **wan**, or **both**. The prefix is then installed on the subscriber interface where the subscriber is instantiated. The prefix state is tied to the state of the subscriber. Once the subscriber session ends, the prefix is removed from the subscriber interface and subsequently from both the FDB and the RIB. This feature can be used as an alternative to unnumbered subscriber interfaces, where the subscriber interface prefix does not need to be predetermined. However, by installing the prefix after authentication, the subscriber interface becomes numbered. In an unnumbered subscriber interface all subscriber routes are installed whereas in a numbered subscriber interface only the subscriber interface prefix is advertised, therefore

reducing the number of advertised routes significantly. The RADIUS-installed prefix can then be advertised through a routing protocol. Subscriber interface prefixes are under the protocol **direct** type similar to other router interfaces. To advertise only the subscriber interface prefix installed by RADIUS, origin **aaa** can be used in the router policy.

The IPv6 node under the group interface contains the DHCPv6 proxy configuration and the router advertisement configuration.

9.3.2.3 64-bit and 128-bit WAN Mode

Subscriber interfaces are created as 64-bit WAN mode interfaces by default. At the time of creation, the subscriber interface can also be created as a 128-bit WAN mode interface. After the subscriber interface is created, the WAN mode cannot be changed. To change the WAN mode, the 64-bit subscriber interface must be removed and then recreated as 128-bit. This section describes the differences between 64-bit and 128-bit WAN modes.

In a 64-bit WAN mode subscriber interface, the following rules apply.

- The system differentiates each subscriber using only the first 64 bits of the WAN address (each host must have a unique /64 prefix, with the exception of bridge host.) This differentiation includes the ability to identify individual subscribers and to apply different subscriber profile and SLA-profiles to each subscriber.
- For IPoE bridge hosts, when a group of hosts shares a prefix, all hosts must share the same SLA-profile.
- Each SLAAC subscriber must use a unique /64 prefix. IPoE bridge hosts can share the same SLAAC prefix and must also share the same SLA-profile.
- Each IPv6 data-trigger host must use a unique /64 prefix.
- The DHCPv6 server must be set up to assign each host with a unique /64 prefix. The 7750 SR local DHCP server assigns each subscriber with a unique /64 prefix by default.

64-bit WAN mode is applicable in deployment models where each subscriber is assigned a unique /64 WAN-prefix which can be used for DHCP or SLAAC.

In a 128-bit WAN mode subscriber interface, the following rules apply.

- It is not recommended to change a subscriber interface from unnumbered to numbered (or vice versa). If the subscriber interface must be changed, all ESM hosts under the subscriber should be removed first.

- The system can uniquely identify each subscriber using the full 128-bit WAN address. Each 128-bit WAN host can have its own unique SLA and Subscriber profile.
- When provisioning a numbered subscriber interface, an IPv6 address or a prefix can be assigned. The mask for the address can range from 32 to 127. If the mask is less than 96, an internal /96 route is generated when a WAN host is created (via DHCP IPv6 IANA). This automatically-generated /96 route is used for subscriber lookup. These /96 routes are hidden and do not appear on the RIB, but do occupy a route entry on the system forwarding table. A /96 route can serve approximately 4.2 billion WAN hosts. Therefore, a single /96 prefix or address should be able to accommodate all WAN hosts terminating on a subscriber interface. Nokia recommends, when using 128-bit WAN mode to configure subscriber interface, use addresses or prefixes with a mask length of 96 to 127. When using 128-bit WAN mode, it is not recommended to assign individual subscribers unique /64 prefixes, because the system generates an internal /96 route for each host, thus overloading the routing table.
- Adding and removing a prefix or address from the subscriber interface in 128-bit WAN mode may trigger /96 routes to be generated or deleted, which can impact subscriber service. Nokia recommends performing this action during off-peak hours.
- 128-bit WAN mode is supported in unnumbered subscriber interfaces. Each WAN host will generate a /128 route.
- The **allow-unmatching-prefixes** command can be performed on a numbered subscriber interface in 128-bit WAN mode. This functionality can be used as a subnet migration tool, but must be performed without hosts under the subscriber -interface. Changing a subscriber interface from numbered to unnumbered (or vice versa) will impact subscriber service.
- IPv6 DHCP IANA subscribers can be assigned incremental 128-bit addresses.
- IPoE-bridge mode is only recommended to be configured with 64-bit WAN mode. For 128-bit WAN mode, the system generates at least one /96 prefix per subscriber to help lookup. Because each subscriber interface has a limited number of allowed prefixes, generating a /96 per subscriber reduces scalability. If 128-bit WAN mode is used, each DHCP IANA host can have a distinct SLA profile. In 64-bit WAN mode, all DHCP IANA hosts must share the same SLA profile.
- For IPoE bridge SLAAC hosts, hosts sharing the same /64 prefix must share the same SLA profile. IPoE bridge SLAAC hosts do not differ from 128-bit or 64-bit WAN mode.
- Each IPv6 data-trigger host can use a unique /128 address.
- The DHCP IA_NA host for both PPPoE and IPoE hosts can assign incremental 128-bit addresses.

- For retail VPRN that requires 128-bit WAN mode support, the wholesale subscriber interface must also be configured with 128-bit WAN mode.
- SLAAC hosts and DHCP WAN hosts must not share the same prefix.
- The following host types are not supported:
  - GTP hosts
  - hybrid access hosts
  - WLAN hosts
  - default hosts

9.3.2.3.1 Migration from 64-bit to 128-bit WAN Mode

It might be beneficial in some deployments for operators to migrate from 64-bit to 128-bit WAN mode. For example, the ability to assign consecutive 128-bit address and minimize the subnet required for Residential Gateway or Cable Modem IPv6 DHCP IANA WAN management address.

A 64-bit WAN mode subscriber interface cannot be changed into a 128-bit WAN mode subscriber interface in real time. To migrate to a 128-bit WAN mode subscriber interface, the 64-bit WAN mode subscriber interface must be a removed and re-created. The 64-bit configuration must be copied, shutdown, and the configuration removed. The configuration can be pasted back with the 128-bit mode added to the subscriber interface. Below are some migration scenarios.

Migration of PPPoE and IPoE DHCP Hosts on MSAPs

Change RADIUS, Diameter, and LUDB in advance of migrations to minimize service impact. Ensure that MSAP stickiness is disabled and idle sticky MSAPs are removed. Nokia recommends performing this migration during a maintenance window.

To prepare to migrate PPPoE and IPoE DHCP hosts on MSAPs, perform the following steps.

1. Create new subscriber and group interfaces for the 128-bit WAN mode.
2. Update the database for the host-related MSAP parameters. This includes AAA (both RADIUS and Diameter) and LUDB. The updated DB directs subscribers to the new subscriber and group interface.

A migration can be performed for either PPPoE and LNS hosts or IPoE DHCP-based hosts. The migration is dependent on subscriber deletion.

For PPPoE and LNS hosts, when a host disconnects their session, the next session is migrated. To speed up the migration, and depending on the RG capability, manually clearing the session could trigger the RG to re-connect via PPPoE immediately, and migrate to the new interface.

When migrating IPoE DHCP-based hosts, Nokia recommends changing both the current DHCPv4 and DHCPv6 lease time and rebind times to one hour or more. It is important to migrate only a small sample size to control the number of DHCP renews. Subscribers are migrated in the following three ways.

- Subsets of subscribers are migrated to the new 128-bit interface without any end user action. For example, the end user does not need to reset their modem or RG. The new authentication establishes the host on the new interface. A maintenance window is not required.
- To migrate the remaining subscribers, enable the drain function on the DHCP server to stop all DHCP lease renewal. However, the DHCPv4 and DHCPv6 lease for a subscriber might end at different times. When both the DHCPv4 and DHCPv6 leases end, the subscriber is removed from the system. Depending on the RG/CM capability, it may send a DHCP request immediately for a new DHCPv4 and DHCPv6 addresses. The subscriber is now created on the new 128-bit subscriber interface without any action by the end user.
- For any remaining subscribers that have not been migrated, after the migration window has waited for the one hour lease time, action by the end user may be required. The operator must shut down both the subscriber interface and group interface, and clear all remaining hosts from those interfaces. The end user is then required to manually reboot the RG to send a DHCP discover/solicit.

Migration of PPPoE and IPoE DHCP Hosts on Static SAPs

Nokia recommends performing this migration during a maintenance window. This migration process is service-impacting.

When migrating IPoE DHCP-based hosts, Nokia recommends changing both the current DHCPv4 and DHCPv6 lease time and rebind time to one hour or more. It is also recommended to migrate only a small sample size to control the number of DHCP renews. After all leases have been changed to a shorter lease time, perform the following steps to prepare for the migration.

1. Remove the old subscriber and group interfaces. This may require manual clearing of subscriber sessions.
2. Re-create new subscriber and group interfaces for the 128-bit hosts (including SAPs).

3. Apply the appropriate AAA (RADIUS and Diameter) and LUDB changes (new 128-bit IPv6 addresses for all hosts), if any.

Following these preparation steps, a migration can be performed for either PPPoE and LNS hosts or IPoE DHCP-based hosts. The migration is dependent on subscriber deletion.

For PPPoE and LNS, when hosts disconnect their session, the RG may try to re-connect via PPPoE immediately and migrate to the new interface.

For IPoE hosts, new subscribers are automatically migrated upon logging in. Some end customers may be required to manually reboot the RG to send a DHCP discover/solicit.

Migration of Data-Trigger Hosts

Nokia recommends performing this migration during a maintenance window. This migration process is service-impacting. Before the migration can begin, the data-trigger node on the group interface must be shut down, then all the data-triggered hosts must be cleared.

To migrate data-trigger hosts:

1. Remove the existing subscriber interface and recreate the new 128-bit subscriber interface, group interface, and related parameters. Removing the existing subscriber interface may require clearing data-triggered subscribers from the system.
2. Update the AAA (RADIUS or Diameter) or LUDB parameters related to the MSAP or SAP. If the end-subscriber is assigned a new IP address, all traffic using the old IP-address is dropped until the new IP address is in use.
3. To complete the migration, the subscriber must send a data packet for authentication using the new assigned IP address.

9.3.2.4 Behavior

9.3.2.4.1 Dual Stack

Clients may support both IPv4 and IPv6 simultaneously (dual stack hosts.) In this case one subscriber host entry will be created for the IPv4 address family and one for the IPv6 instance. The scaling limits apply for all entries, regardless of address type.

For DHCP, these subscriber hosts are fully independent (as they are set up through different protocols), but for PPPoE hosts or RGs, the ESM information in both subscriber host entries is linked together through the PPPoE session.

9.3.2.4.2 Router Advertisements (RA)

Router Advertisement (RA) messages begin immediately after the subscriber host is instantiated and unsolicited messages are sent in the interval defined in the configuration. Apart from unsolicited RAs, the client may also send a router solicitation (RS) to explicitly request the information. RAs are throttled so that they are not sent more than once every three seconds.

The Router Advertisement Policy feature overrides the group interface RA configuration for hosts on a given MAC on a given SAP. The policy is applied directly to the sending instance where it sends periodic RAs. The policy can be applied at authentication or by CoA during the subscriber session. The RA policy can be used in the following ways.

- When applied to an IPoE session or a PPPoE session, the RA policy is applied to all IPv6 hosts within the session.
- When applied to a subscriber (regardless of whether it is session-based), the RA policy is applied to all IPv6 hosts within the subscriber.
- When applied to a host within a subscriber, the RA policy is applied only to the IPv6 hosts for the particular MAC (for IPoE session) or the particular MAC and PPP session (for PPPoE session).

The prefix option inside the RA policy allows independent prefix options for subscribers that utilize bridge hosts. The bridge hosts can consist of both DHCPv6 and SLAAC, and are represented as stateful and stateless within the policy respectively. Within the policy, the **autoconfig** flag is not configurable and is disabled by default for the DHCPv6 address and enabled by default for SLAAC. For SLAAC hosts, if the **autoconfig** flag is enabled inside the RA policy along with the SLAAC prefix, the **autoconfig** flag for the DHCPv6 address or prefix is not enabled as a result. The timers for either SLAAC and DHCPv6 prefixes can also be configured independently.

The router advertisement policy has a separate configuration for stateless and stateful operations. The general recommendation is to configure the valid and preferred lifetimes for longer than the minimum RA interval to ensure the subscriber has a valid address to use between each RA interval. If this general rule is not followed, the subscriber can deprecate the SLAAC prefix between each RA interval and experience service interruptions. As the minimum RA interval is approximately 15 minutes, the valid and preferred lifetime values should be at least 15 minutes. Shorter valid and preferred lifetime values can impact the system’s scalability. The

stateful RA has a static option and a dynamic option when configuring the valid and preferred lifetime values. If the static option is used, the valid and preferred lifetime values should be greater than the RA interval. For the dynamic option, the **auto-lifetimes** feature derives the valid and preferred lifetime values from the DHCPv6 lease. Therefore, the RA and DHCPv6 will have the same valid and preferred lifetime values.

SLAAC hosts are assigned prefixes, where the full Global Unicast Address (GUA) is not known. Regardless of the **force-mcast** configuration, the destination IP address for an RA to an SLAAC host will always be a multicast IP address, with one exception. If the feature **allow-multiple-wan-address** is enabled and the same host (same MAC on the same SAP and same device) has a DHCPv6 NA address, the NA address will be used for the unicast RA. The MAC address can either be a multicast or unicast address, depending on the configuration of **force-mcast**.

Table 75 outlines the behavior of the system when the RA policy VSA is included in authentication, CoA, and re-authentication. The RA policy that is sent from RADIUS may not yet be provisioned in CLI, and therefore may not exist in the system.

Table 75 **RA Policy Behavior**

	Authentication	**CoA/Tools CoA**	**Re-authentication**
BRG	An RA policy does not need to exist. The RA policy becomes active when a matching RA policy is provisioned. If an RA policy does not exist, the RA parameters configured under the group interface are used.	An RA policy must exist; otherwise, a NACK is sent in response to the CoA.	An RA policy must exist; otherwise, all VSAs and the RA policy are ignored. An SNMP trap is raised.
Subscriber is session-based (for example, an IPoE session)	An RA policy does not need to exist for the IPv4 host. The RA policy becomes active when a matching RA policy is provisioned. ESM IPv6 host creation fails when a policy does not exist. If an RA policy does not exist, the RA parameters configured under the group interface are used.	An RA policy must exist; otherwise, a NACK is sent in response to the CoA.	An RA policy must exist; otherwise, all VSAs and the RA policy are ignored. An SNMP trap is raised.

Table 75 **RA Policy Behavior (Continued)**

	Authentication	CoA/Tools CoA	Re-authentication
Subscriber has a dual stack host and is not session-based	An RA policy does not need to exist for the IPv4 host. The RA policy becomes active when a matching RA policy is provisioned. ESM IPv6 host creation fails when a policy does not exist. If an RA policy does not exist, the RA parameters configured under the group interface are used.	An RA policy must exist; otherwise, a NACK is sent in response to the CoA.	An RA policy must exist; otherwise, all VSAs and the RA policy are ignored. An SNMP trap is raised.
IPv4 host that is not session-based	An RA policy must exist. Otherwise, the subscriber setup is rejected.	An RA policy must exist; otherwise, a NACK is sent in response to the CoA.	An RA policy must exist; otherwise, all VSAs and the RA policy are ignored. An SNMP trap is raised.
Dual stack host that is not session-based, where the CoA is targeted to an IPv4 host only	n/a	An RA policy must exist; otherwise, a NACK is sent in response to the CoA.	n/a
Dual stack host that is not session-based, where the CoA is targeted to an IPv6 host only	n/a	An RA policy must exist; otherwise, a NACK is sent in response to the CoA.	n/a
IPoE linking (both session-based and not session-based)	An RA policy does not need to exist for the IPv4 host. The RA policy becomes active when a matching RA policy is provisioned. ESM IPv6 host creation fails when a policy does not exist. If an RA policy does not exist, the RA parameters configured under the group interface are used.	An RA policy must exist; otherwise, a NACK is sent in response to the CoA.	An RA policy must exist; otherwise, all VSAs and the RA policy are ignored. An SNMP trap is raised.

Table 75 **RA Policy Behavior (Continued)**

	Authentication	CoA/Tools CoA	Re-authentication
PD host as managed to IPv4 (both session-based and not session-based)	An RA policy does not need to exist for the IPv4 host. The RA policy becomes active when a matching RA policy is provisioned. ESM IPv6 host creation fails when a policy does not exist. If an RA policy does not exist, the RA parameters configured under the group interface are used.	An RA policy must exist; otherwise, a NACK is sent in response to the CoA.	An RA policy must exist; otherwise, all VSAs and the RA policy are ignored. An SNMP trap is raised.

Router Advertisement Policy Limitations

The following are RA policy limitations.

- As a result of a general limitation on CoA, a maximum of 32 bridge hosts can be updated. This limitation exists in the case of BRG.
- RA policies are not supported for DSM.
- RAs are configured to be sent at certain intervals. It is highly recommended that this interval is considered when configuring the prefix lifetimes. For example, if the interval is configured for one hour and the prefix lifetime is configured for 30 minutes, the (SLAAC) host will be removed from the BNG before the next RA is sent.
- If the parameters within the RA policy are modified, the parameters only take effect on the next interval.

9.3.2.4.3 CoA and Disconnect-Request

For IPv6 subscriber hosts, RADIUS-triggered mid-session changes and session terminations may identify the subscriber host to be changed by the same address or prefix that was originally returned from RADIUS. Only one address attribute (framed-IP address, framed-IPv6-prefix, delegated-IPv6-prefix or Alc-IPv6-address) may be given in a single request.

For PPPoE clients, changing either the IPv4 or IPv6 information will result in both the v4 and v6 subscriber host being modified (if they are contained within the same PPPoE session.)

The only CoA action that is allowed for IPv6 hosts is a change of ESM strings; creation of new hosts and forcing a DHCPv6 RENEW is not supported.

9.3.2.5 Delegated-Prefix-Length

The delegated prefix length (DPL) is applicable to subscriber-hosts with IPv6 Prefix (IA-PD) assigned by the DHCPv6 Server. IPv6 Prefix is more akin to a route then it is to an IP address. The length of the prefix plays crucial role in forwarding decisions, antispoofing, and prefix assignment via DHCPv6 pools in the local DHCPv6 Server.

The structure of an IPv6 prefix is shown in Figure 70.

Figure 70 **IPv6 Prefix**

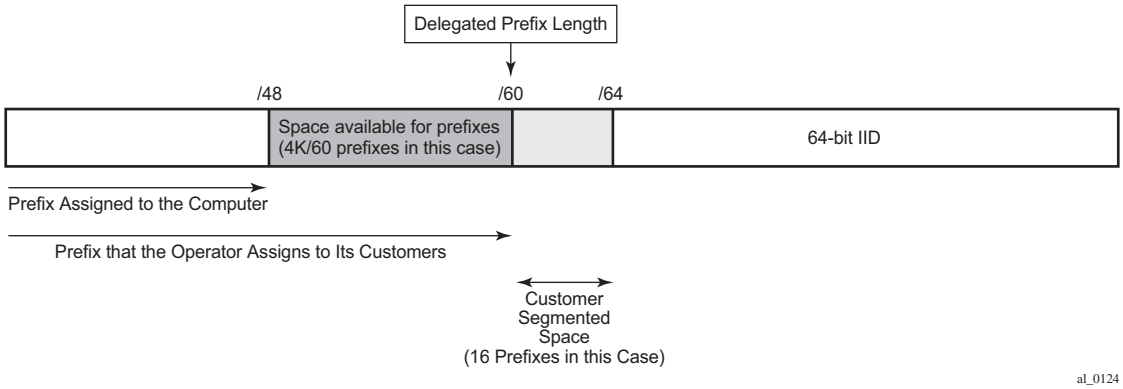

For example, a DHCPv6 server prefix pool contains an aggregated (configured) IPv6 prefix from which the delegated prefixes will be carved out. In Figure 70 this aggregated IPv6 prefix has length of /48. In addition, the DHCPv6 server needs to know the length of the delegated prefix (in the above case /60). These two values are marking the boundary within which a unique delegated prefix will be selected. This is represented by the purple area in Figure 70.

The delegated prefix length can be obtained using:

- RADIUS
  - Delegated-IPv6-Prefix attribute that contains the prefix and the length (Delegated-IPv6-Prefix = AAAA:BBBB::/56). The DPL in this case is /56.
  - Alc-Delegated-IPv6-Prefix-Length VSA (to be used in conjunction with the DHCPv6 pool name - Alc-Delegated-IPv6-Pool VSA)
- LUDB – configured via LUDB per IPoEv6/PPPoEv6 host:

  This is to be used along with the DHCPv6 pool name (ipv6-delegated-prefix-pool) defined under the same CLI hierarchy.

```
configure
subscriber-mgmt
     local-user-db <name>
          dhcp | ppp
               host <name>
                    ipv6-delegated-prefix-length
                      [48 to 64]
```

Alternatively, the entire prefix, including the DPL can be returned via LUDB.

```
configure
     subscriber-mgmt
          local-user-db <name>
                    dhcp | ppp
                         host <name>
                              ipv6-delegated-prefix
                                <ipv6-prefix/prefix-
                                length>
```

- DHCPv6 server – each DHCPv6 pool can optionally be configured with a DPL:

```
configure
      service/router
          dhcp6
               local-dhcp-server <name>
                    pool <pool-name>
                         delegated-prefix-length [48 to
                           127]
```

- Configured statically under the **ipv6** CLI node of **subscriber-interface**. In this case, the DPL is fixed for all subscriber hosts under the subscriber interface.

```
configure
     service ies/vprn
          subscriber-interface <name>
                    ipv6
                         delegated-prefix-length [48 to
                           64] | variable
```

9.3.2.5.1 Order of Preference for DPL

In case that the DPL is statically provisioned under the subscriber-interface>ipv6> hierarchy, all hosts under this subscriber-interface will inherit this fixed DPL. In case that the DPL is provided via LUDB or RADIUS in addition to static configuration under the subscriber-interface then the LUDB or the RADIUS one not match the DPL that is statically provisioned under the subscriber-interface. Otherwise, the prefix instantiation in 7450 ESS and 7750 SR will fail.

Note that the **no delegated-prefix-length** command under the **subscriber-interface>ipv6>** hierarchy means that the DPL is set to a default-value of 64.

When the delegated-prefix-length commands under the **subscriber-interface>ipv6** hierarchy is set to variable, prefixes under such subscriber-interface can have different lengths and the DPL can be configured via one of the following means:

- LUDB
- RADIUS
- DHCP Server

9.3.2.5.2 DHCP Server Address Utilization and Delegated Prefix Length

In case that the delegated prefix length is variable, for each consecutive address allocation request for the given delegated prefix, the DHCPv6 server will allocate the prefix at the end of the last delegated lease with the same delegated prefix length. This will minimize the address space fragmentation within the configured prefix.

9.3.2.6 DHCPv6 Relay Agent

A DHCPv6 Relay Agent can support a 7450 ESS and 7750 SR DHCPv6 local server (same or remote chassis) and a third party DHCPv6 external server.

An incoming DHCPv6 client message is relayed within the Relay-Forward message specified in RFC 3315. If the server responds with a valid address/prefix, the ESM process attempts to install it. If it fails, the DHCPv6 Relay Agent sends an explicit RELEASE to the server. There is no retransmission of DHCPv6 Relay-Forwards in the case of failure – it requires the client to re-start or re-send the original DHCPv6 message.

A Lightweight DHCPv6 Relay Agent may insert Relay Agent Information including the Interface ID option between the DHCPv6 client and the DHCPv6 Relay Agent.

Additional Relay Agents (non-LDRA) between the DHCPv6 client and the DHCPv6 Relay Agent are not supported.

DHCPv6 Reconfigure messages received from an external DHCPv6 server are forwarded to the DHCP client, if a corresponding DHCPv6 lease exists. The Reconfigure message can be sent in a unicast message to the client or encapsulated in a Relay-Reply message to the DHCPv6 relay agent. The DHCPv6 Reconfigure message is dropped if no corresponding DHCPv6 lease exists.

9.3.2.6.1 Configuring a DHCPv6 Relay Agent

A DHCPv6 Relay Agent is configured in the IPv6 DHCP6 context of a group-interface:

```
config>service>vprn>sub-if>grp-if>ipv6>dhcp6# relay ?
config>service>ies>sub-if>grp-if>ipv6>dhcp6# relay ?
  - no relay
  - relay

 [no] client-applications - Configure the set of DHCP6 relay server client
                            applications
 [no] description     - Description for DHCPv6 relay
 [no] link-address    - Configure the link address of the DHCPv6 relay messages
 [no] option          + Configure the DHCPv6 Relay information options
 [no] server          - Configure the DHCPv6 server IPv6 address
 [no] shutdown        - Administratively enable/
disable DHCPv6 relay on this interface
 [no] source-address  -
 Configure the source IPv6 address of the DHCPv6 relay messages
```

Up to eight DHCPv6 servers can be provisioned to be served by a DHCPv6 Relay Agent. A Relay-Forward is send to all servers and the Relay-Replies from all servers are sent to the client.

The “client-applications” parameter specifies if the Relay Agent can be used for IPoE (dhcp) or PPP (ppp) hosts.
Optional configuration parameters:

- description — A free configurable description string.
- link-address — The link address field in the DHCPv6 Relay-Forward message header.

  The link address can be configured to enable link-address based pool selection in a 7450 ESS and 7750 SR DHCPv6 local server. The address must be one of the IPv6 prefixes configured at the ipv6 subscriber-prefixes context for a subscriber interface. If not configured, the system selects one of the prefixes.
- option: allows to configure following options to be inserted in the Relay-Forward message:
  - Interface-Id [18] — The interface ID option identifies the interface on which the DHCPv6 client message is received. The format options are the following:
    - ascii-tuple: — *host-name*|*service-id*|*group-interface-name*|*sap-id*
    - ifindex — Interface index for the group-interface
    - sap-id — SAP identifier (port and VLANs)
    - string <string>: — A free configurable string (max. 80 chars)

- Remote-Id [37] — Relay Agent Remote Id option contains the DHCPv6 client DHCP Unique Identifier (DUID).

- source-address: the source-address of the Relay-Forward messages.

  If not configured, the outgoing interface IPv6 address is used. The source-address configuration is mandatory for a DHCP Relay Agent in a VPRN service when the DHCPv6 server is reachable via a tunneled next-hop (MPLS).

9.3.2.7 DHCPv6 Relay to Third Party DHCPv6 External Server

When the DHCPv6 Relay Agent is relaying to a third party DHCPv6 external server, following conditions should be met:

- The third party DHCPv6 server must return a unique IA_PD IPv6 delegated prefix (/64 or lower) for each allocation. The length of the IA_PD IPv6 delegated prefix must match the delegated-prefix-len configured on the subscriber interface on the 7750 DHCP L3 relay. This length is also included in the Relay-Forward message as PFX_LEN option (3) in a Vendor-Specific-Information-Option (17).
- For IPv6oE routed CPE's, the 3rd party DHCPv6 server must return a unique IA_NA IPv6 address (/128) from a different /64 subnet for each allocation.
- For IPv6oE hosts behind bridged CPE's,
  - the third party DHCPv6 server must return a unique IA_NA IPv6 address (/128) from a different /64 subnet for each allocation (host) that belongs to a different CPE.
  - the third party DHCPv6 server may return a unique IA_NA IPv6 address (/128) from the same /64 subnet for allocations (hosts) that belong to the same CPE and that are attached to the same vlan (SAP) on the BNG.

Following information is available to the third party DHCPv6 server in a Vendor-Specific-Information-Option (17) included in the Relay-Forward message:

- WAN_POOL option (1) contains the pool name from which the IA_NA IPv6 address should be allocated.
- PFX_POOL option (2) contains the pool name from which the IA_PD IPv6 delegated prefix should be allocated.
- PFX_LEN option (3): contains the IA_PD IPv6 delegated prefix length that should be allocated.

9.3.2.8 DHCPv6 Local Server

A local DHCPv6 pool server for both addresses (IA_NA) and prefixed (IA_PD) manages the address and prefixes sent to either routing gateways or hosts.

Because IPv6 home networks lack NAT, the IPv6 addresses delegated to a routing gateway are in turn assigned to hosts in the home. These addresses are assigned with reasonably long (but configurable) lifetimes such that the loss of the WAN connection will not result in the IPv6 hosts in the LAN losing their IPv6 addresses. One consequence of these long lifetimes is that the IPv6 hosts will retain any IPv6 address provided the valid-lifetime is greater than zero. Should an operator delegate a prefix and then at a later time delegate a second IPv6 prefix, a host may end up with two or more valid prefixes. This situation plays havoc with IPv6 source address selection and may result in impaired service.

To overcome the problems of multiple IPv6 prefixes in the home, the operator must ensure that the individual subscriber has the same IPv6 prefix even across modem reboots (that is, if a subscriber session is destroyed and later re-created, an attempt should be made to use the previously delegated prefix). In release 8.0, the operator used RADIUS for all address and prefix assignment, but in release 9.0, with the introduction of the local DHCPv6 server, it requires the 7750 to process and maintain some state even after a session disconnects.

For the DHCPv6 local server to function, a DHCPv6 relay or proxy function must also operate alongside ESM. For the purposes of this document, to relay means to implement a DHCPv6 Relay as indicated in RFC 3315 : a relay encapsulates the client DHCP message within a DHCP Relay-Forward message and unicasts it to a specified destination.

A proxy is an internal concept. Unlike a DHCPv6 relay, the DHCPv6 proxy does not encapsulate the client message in a Relay-Forward, nor does it send packets towards the Local DHCPv6 Server. The DHCPv6 proxy is exclusively used as an interface between the RADIUS Access-Accept or local user database lookup and the DHCPv6 client in the consumer device.

The use of the DHCPv6 relay or proxy function depends on the attributes returned from authentication phase (RADIUS or LUDB).

1. DHCPv6 Proxy:
   - If only IPv6 address/prefix information provided (Framed-IPv6-Prefix, Alc-IPv6-Address or Delegated-IPv6-Prefix).
2. DHCPv6 Relay:
   - If no IPv6 address/prefix (Framed-IPv6-Prefix, Alc-IPv6-Address or Delegated-IPv6-Prefix) and no IPv6 pool (Framed-Pool, Delegated-Pool) information provided.

– If no IPv6 address/prefix (Framed-IPv6-Prefix, Alc-IPv6-Address or Delegated-IPv6-Prefix) and IPv6 pool (Framed-Pool, Delegated-Pool) information provided.

3. If both IPv6 address/prefix (Framed-IPv6-Prefix, Alc-IPv6-Address or Delegated-IPv6-Prefix) and IPv6 pool (Framed-Pool, Delegated-Pool) information are present, the DHCP packet is DROPPED.

Note: If IPv6 DNS parameters are returned in RADIUS and a pool is specified then the DNS parameters are ignored. It is the DHCPv6 server that will need to reply with appropriate DNS servers.

9.3.3 Dynamic Subscriber Host Processing

9.3.3.1 Dynamic Tables

To support all processing for Enhanced Subscriber Management, several tables are maintained in the router (Figure 71).

Figure 71 **Enhanced Subscriber Management Dynamic Tables**

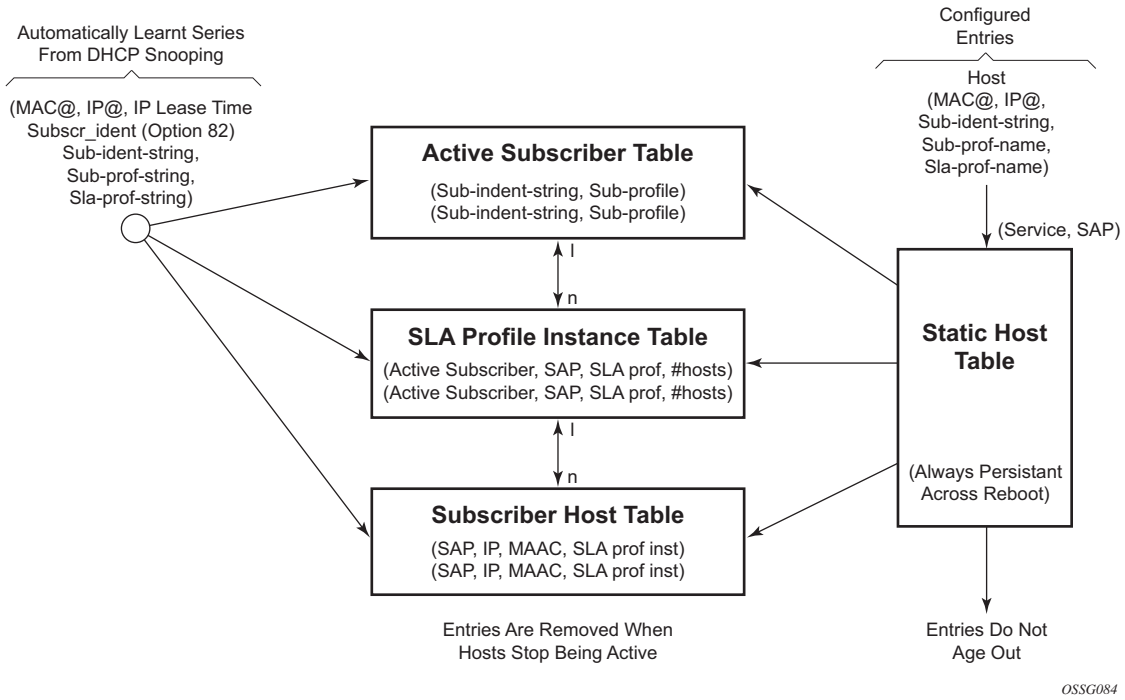

9.3.3.1.1 Active Subscriber Table

An entry is created in the active subscriber table when the first host (either dynamic or static) is created with a certain subscriber identification string. The entries are grouped by their subscriber identification string.

Fields for each entry in the active subscriber table include:

- The subscriber identification string (see Subscriber Identification String).
- In use subscriber profiles (see Subscriber Profile).

9.3.3.1.2 SLA Profile Instance Table

An entry is created in the SLA profile instance table when the first subscriber host on a certain SAP is created that uses a certain SLA profile. All subsequent hosts of the same subscriber on the same SAP that use the same SLA profile will be associated with this entry. When the last host on this SAP, using this SLA profile disappears, the SLA profile instance is deleted from the table and the associated queues are removed.

SLA profile instances can not span multiple subscriber SAPs. If subscriber hosts from the same subscriber exist on multiple SAPs and are associated with the same SLA profile template, a separate SLA profile instance is created for each SAP.

Fields for each entry in the SLA profile instance table include:

- Active subscriber
- SAP
- SLA profile
- Number of active subscriber hosts that share this instance

9.3.3.1.3 Subscriber Host Table

An entry is created in the subscriber host table if anti-spoofing is enabled as well as:

- The first host (dynamic or static) with a specific IP and MAC combination is created. If the anti-spoof is IP only, the MAC address is masked to all 0's. If anti-spoof is MAC, only the IP address is 0.0.0.0. All dynamic hosts and static hosts with the same IP and MAC combination will be associated with the same subscriber host entry. If the anti-spoof type includes IP (IP-only or IP/MAC), there can be at most two hosts associated with the entry: one dynamic and one static. If the anti-spoof type is MAC-only, there can be a combination of several dynamic and static hosts associated with the entry.
- The non-prof-traffic is provisioned. Both IP and MAC address are all 0's.

Fields for each entry in the subscriber host table include:

- SAP
- IP address
- MAC address
- SLA profile instance (enhanced mode only)

9.3.3.1.4 DHCP Lease State Table

An entry in the DHCP lease state table is created for each dynamic host. Fields for each entry in the lease state table include:

- Assigned IP address
- Assigned MAC address
- Persistence key

9.3.4 Enhanced Subscriber Management Entities

Figure 72 illustrates the relationship between the main entities in Enhanced Subscriber Management:

- A subscriber is associated with only one subscriber profile.
- A subscriber can be associated with one or more SLA profile (a VPLS service with 2 different SAPs can have different SLA profiles for the same subscriber).
- A maximum of one SLA profile instance is generated (including ingress and egress queues) per SAP per SLA profile.
- One or more hosts can be assigned to each SLA profile instance (these will share the same queues).

Figure 72 **Relationship Between Enhanced Subscriber Management Entities**

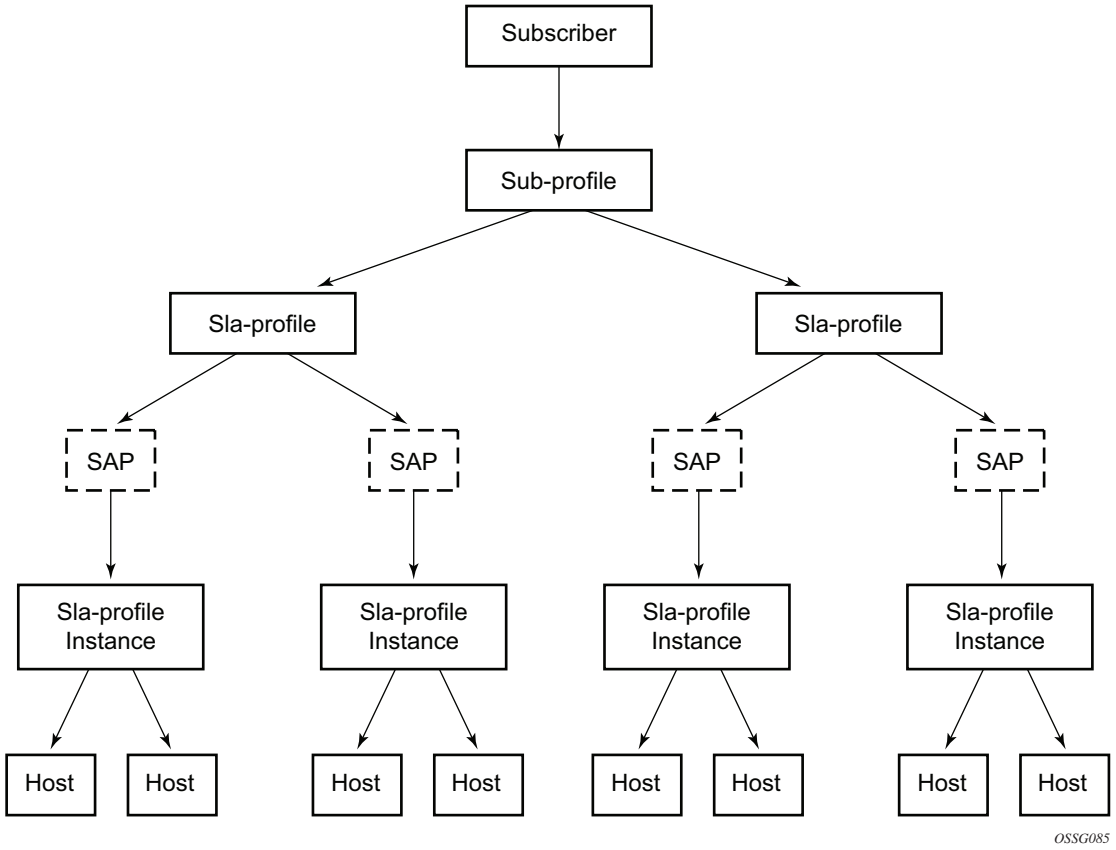

9.3.4.1 Instantiating a New Host

When a DHCP ACK is received for a new subscriber host on a particular SAP:

- The ACK message is parsed using the appropriate script.
- An entry is generated in the subscriber host table with indices:
  - The SAP on which the host resides
  - The assigned IP address
  - The assigned MAC address and as lookup parameters:
    - The subscriber profile and
    - The SLA profile to be used (derived from using the script).

If this is the first host of a subscriber, an HQoS scheduler is instantiated using the ingress and egress scheduler policies referred to in the subscriber profile. Otherwise, if the subscriber profile of the new host equals the subscriber profile of the existing subscriber, the new host is linked to the existing scheduler. If the subscriber profile is different from the subscriber profile of the existing subscriber, a new scheduler is created and all the hosts belonging to that subscriber are linked to this new scheduler. Notice that the new subscriber profile will not conflict with the subscriber profile provisioned for a static host or non-sub-traffic under the same SAP.

If this is the first host of a subscriber on a particular SAP using a particular SLA profile, an SLA profile instance is generated and added to the SLA profile instance table. This includes instantiating a number of queues, according to the ingress and egress QoS profiles referred to in SLA profile, optionally with some specific overrides defined in the SLA profile. Otherwise the host is linked to the existing SLA profile instance for this subscriber on this SAP.

Note: Any QoS and IP filter policies defined on the SAP are still processed even if Enhanced Subscriber Management is enabled on the SAP. For IPv4 traffic that is dropped due to anti-spoofing, counters, logging, and mirroring can be used. All other Layer 2 traffic that is never blocked by anti-spoofing can be processed by applying a QoS policy on the SAP and can still be classified differently, by the dot1p value.

Note: If insufficient hardware resources (queues) or software resources (profile instances) are available to support the new host, the DHCP ACK is dropped and an event is generated.

9.3.4.2 Packet Processing for an Existing Host

Whenever an IP packet arrives on a subscriber-facing SAP on which Enhanced Subscriber Management (ESM) is enabled, a lookup is done in the subscriber host table using as the index the SAP, source IP address, and source MAC address.

- If there is no entry, this means that the host is not using his assigned IP address, so the packet is dropped;
- If there is an entry, this will refer to the subscriber profile and SLA profile to be used.

9.3.5 ESM Host Lockout

This feature increasingly penalizes hosts that fail repeated login attempts within a configurable time interval. This is done by holding off on creation attempts for these hosts for a configured but adaptable time period. A transient failure, due to a mis-configuration, is quickly corrected and does not prevent the host from logging in within a reasonable amount of time. At the same time, a malicious client or a constantly mis-configured client is locked-out and will not take up resources impacting other clients.

A lockout time per host supports exponential back-off with each retry and failure cycle, starting with a configured minimum value and increasing up to a configured maximum. The lockout time can be reset to the configured minimum value if there is no failed retry within a configured time threshold. The configurable values include:

```
lockout-reset-time seconds
lockout-time [min seconds] [max seconds]
max-lockout-hosts hosts
```

If multiple retries/failure cycles occur within the lockout time, then lockout period is exponentially increased starting from configured minimum value up to the configured maximum value. The lockout is reset to the minimum value if there is no failed retry till this lockout time.

This mechanism is supported for both single and dual-stack PPPoE and IPoE (DHCP) hosts over 1:1 or N:1 static or managed SAPs. The hold-off timer maintenance is on a per host basis (as follows):

- For 1:1 VLAN (PPPoE or IPoE hosts) per <VLAN, MAC address>
- For N:1 VLAN (PPPoE or IPv4oE hosts) per <VLAN, agent-circuit-id, agent-remote-id, MAC@>
- For 1:1 VLAN (IPv6oE hosts) per <VLAN, DUID>

A show lockout state for hosts is supported, given one or more of <SAP, MAC@, agent-circuit-id, agent-remote-id>.

A clear lockout state is supported for hosts given one or more of <SAP, MAC@, agent-circuit-id, agent-remote-id>.

Any changes in configured lockout values will not apply to hosts currently under lockout and will only apply once these hosts are out of lockout.

9.3.5.1 Functionality

ESM lockout is supported for dual-stack PPPoE hosts, L2TP LAC hosts, dual-stack IPoE hosts, and ARP hosts. ESM Lockout tracks the following:

- PPPoE PADI and PADR
- DHCPv4 discover, DHCPv4 request, DHCPv6 solicit, DHCPv6 request
- ARP Request
- PPPoE session disconnect after successful session establishment

During lockout, authentication and ESM host creation is suppressed. A lockout context will be created when a client first enters lockout. The context maintains state and timeout parameters for the lockout. If a lockout policy is configured for the underlying SAP for a host that has failed authentication or host creation, the host enters lockout for the configured minimum time (1 to 86400 seconds). When the lockout time expires, normal authentication and ESM host creation will be resumed on relevant PPP or DHCP messages. In case of another failure, the host will again enter the lockout state. The lockout time for the host on each failure will be exponentially increased up to the configured maximum time (1 to 86400 seconds). The lockout time for a client will be reset to the configured minimum value, and the corresponding lockout context will be deleted, if there is no authentication (and host creation) failure within a configured amount of time that needs to elapse after the client initially enters lockout. This time is called the **lockout-reset-time**.

The host identification for lockout includes <SAP, MAC@, circuit ID, remote ID>.

9.3.6 ANCP and GSMP

9.3.6.1 ANCP

Access Node Control Protocol Management (ANCP) can provide the following information to the router:

- ANCP can communicate the current access line rate to the router. This allows the router to adjust the H-QoS subscriber scheduler with the correct rate or potentially change alarm when the rate goes below a set threshold. This allows a policy manager to change the entire policy when the rate drops below a minimal threshold value. The ANCP actual upstream synchronization rate is mapped to the ingress while ANCP actual downstream synchronization rate is mapped to the egress.

- The router can send DSL line OAM commands to complete an OAM test from a centralized point or when operational boundaries prevent direct access to the DSLAM.

When ANCP is used with Enhanced Subscriber Management (ESM), a new string ancp-string can be returned from the Python script or from RADIUS. If not returned it defaults to the subscriber ID.

ANCP version 0x31 and 0x32 are both supported and will be auto detected at the start of each ANCP session. Within version 0x32, partitioning is also supported.

Multiple partitions from the same Access Node are also supported. If partitions are used, they are automatically detected during the start of an ANCP session.

9.3.6.1.1 Static ANCP Management

As depicted in Figure 73, a DSLAM is connected to an aggregation network that is connecting the DSLAM to a BRAS. ANCP is used to provide SAP level rate management. The DSLAM in this application maintains multiple ANCP connections. The primary connection is to the BRAS, providing rate and OAM capabilities while the secondary is to the router to provide rate management.

7750 SR and 7450 ESS:

Figure 73 **Static ANCP Management Example**

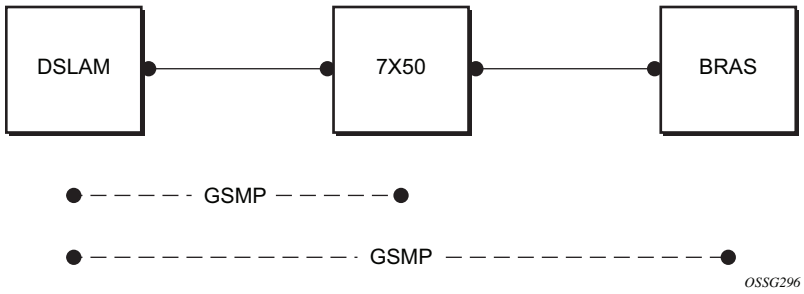

9.3.6.1.2 Enhanced Subscriber Management (ESM) Dynamic ANCP

In this application ANCP is used between the DSLAM and the router to provide line control. There are multiple attributes defined as described below. Figure 74 depicts the connectivity model.

This application is used to communicate the following from the DSLAM to the router (the policy control point):

- Subscriber rate
- OAM

Figure 74 **ESM Dynamic ANCP Example**

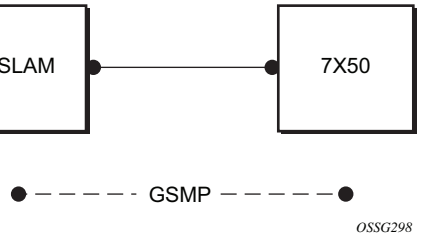

9.3.6.1.3 ANCP String

To support node communication with the access device the line rate, OAM commands, and so on. the node can use an ANCP string that serves as a key in the out-of-band channel with the access node. The string can be either provisioned in the static case, retrieved from RADIUS or from the Python script.

9.3.6.1.4 ANCP Persistency Support

Persistency is available for subscriber's ANCP attributes and is stored on the on-board compact flash card. ANCP data will stay persistence during an ISSU as well as nodal reboots. During recovery, ANCP attributes are first restored fully from the persistence file and incoming ANCP sessions are temporarily on hold. Afterwards new ANCP data can overwrite any existing values. This new data is then stored into the compact flash in preparation for the next event.

9.3.6.2 General Switch Management Protocol Version 3 (GSMPv3)

General Switch Management Protocol version 3 (GSMPv3) is a generic protocol that allows a switch controller node to establish and maintain connections with one or more nodes to exchange operational information. Several extensions to GSMPv3 exist in the context of broadband aggregation. These extensions were proposed to allow GSMPv3 to be used in a broadband environment as additional information is needed to synchronize the control plane between access nodes (such as DSLAMs) and broadband network gateways (such as BRAS).

In the TPSDA framework, nodes fulfill some BRAS functionality, where per subscriber QoS enforcement is one of the most important aspects. To provide accurate per-subscriber QoS enforcement, the network element not only knows about the subscriber profile and its service level agreement but it is aware of the dynamic characteristics of the subscriber access circuit.

The most important parameters in this context are the subscriber-line capacity (DSL sync-rate) and the subscriber's channel viewership status (the actual number of BTV channels received by the given subscriber in any point in time). This information can be then used to adjust parameters of aggregate scheduling policy.

Besides, the above-mentioned information, GSMPv3 can convey OAM information between a switch controller and access switch. The node can operate in two roles:

- As the intermediate controller — The router terminates a connection from the DSLAM.
- As the terminating controller— The router fulfills full the roll of BRAS.

The DSL forum working documents recommends that a dedicated Layer 2 path (such as, a VLAN in an Ethernet aggregation network) is used for this communication to provide a certain level of security. The actual connection between DSLAM and BRAS is established at TCP level, and then individual messages are transported.

9.3.6.3 DHCP Release Messages

The node supports DHCP release messages. A DHCP release message removes state from the DHCP server when the node rejects ACKs or removes hosts.

9.3.6.3.1 DHCP Release

DHCP release messages will be controlled by the node and sent to the DHCP server to clear stale state. There are two examples:

1. If the node drops a DHCP ACK (because of resources, duplicate host or other reasons) the servers state must be cleared and the node will send a DHCP release.
2. When a host state is removed, based on SHCV, ANCP, user clear, and so on, the node will send a DHCP release to the server and the MAC will be flushed from SDPs. A new flag will allow the user to elect not to send the release message If when using a **clear lease** command the host was removed by the user (using a clear command) a new flag will allow the user to elect not to send the release message.

9.3.6.4 DHCP Client Mobility

Client mobility allows the node to use host monitoring (SHCV, ANCP, split DHCP) to remove network and server state when a host is removed locally. This allows for MAC addressed learned and pinned to move based on policy parameters.

Subscriber Host Connectivity Verification (SHCV) configuration is mandatory. This allows clients to move from one SAP to another SAP in the same service. This is only applicable in a VPLS service and group interfaces.

The first DHCP message on the new SAP with same MAC address (and IP address for group-interfaces) will trigger SHCV and will always be discarded.

SHCV will check that the host is no longer present on the SAP where the lease is currently populated to prevent spoofing. When SHCV detects that the host is not present on the original SAP, the lease-state will be removed. The next DHCP message on the new SAP can initiate the host.

9.3.6.5 DHCP Lease Control

DHCP lease control allows the node to be configured to present a different lease to the client. This can be used to monitor the health of the client.

9.3.7 Using Scripts for Dynamic Recognition of Subscribers

Whenever a host belonging to a subscriber is activated (when a PC or set-top box (STB) is turned on), the host will typically request an IP address from the network using DHCP. See the DHCP Management section for an explanation of DHCP and DHCP snooping in the router.

The DHCP ACK response from the DHCP server can be parsed and the contents of the message can be used to identify the class to which this host belongs, and thus, the QoS and security settings to apply.

The information necessary to select these settings can be codified in, the IP address by the DHCP server and/or the Option 82 string inserted by the DSLAM or other access node.

9.3.7.1 Python Language and Programmable Subscriber Configuration Policy (PSCP)

PSCP is an identification mechanism using the Python scripting language. The PSCP references a Python script that can use regular expressions to derive the sub-ident-string, sub-profile-string and sla-profile-string from the DHCP response. A tutorial of regular expressions is beyond the scope of this guide, and can be found on the Internet (refer to http://www.amk.ca/python/howto/regex/).

A tutorial of Python is beyond the scope of this guide but can be found on the Internet (refer to http://www.python.org/).

Example scripts, using some regular expressions, can be found in Sample Python Scripts. See the Python Script Support for ESM section for additional information about the service manager scripting language.

One or more scripts can be written by the operator and stored centrally on a server (in a location accessible by the router). They are loaded into each router at bootup.

Note that if a centrally stored script is changed, it is not automatically re-loaded onto the router. The reload must be forced by executing the **shutdown** and **no shutdown** commands on the affected URL(s).

9.3.7.2 Determining the Subscriber Profile and SLA Profile of a Host

Figure 75 describes the data flow while determining which subscriber profile and SLA profile to use for a certain subscriber host based on a snooped/relayed DHCP ACK for that subscriber host.

Figure 75 **Data Flow in Determining Subscriber Profile and SLA Profile**

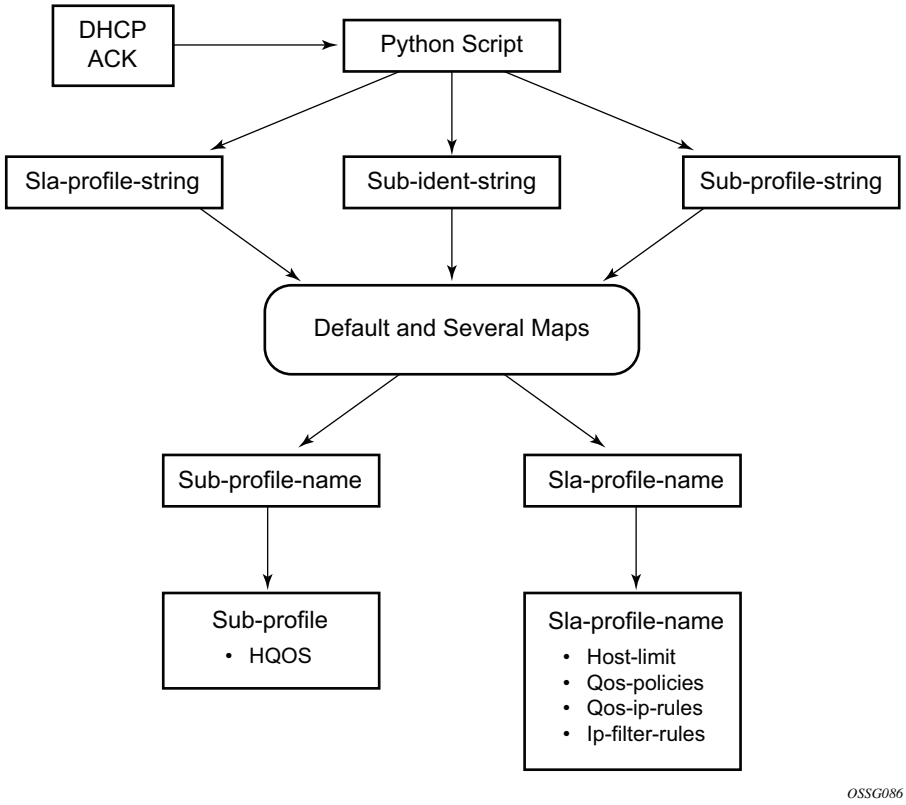

An incoming DHCP ACK (relayed or snooped) is processed by the script provisioned in the sub-ident-policy defined in the SAP on which the message arrived. This script outputs one or more of the following strings:

- sub-ident — Identifies the subscriber (always needed).
- sub-profile — Identifies the subscriber class (optional).
- sla-profile — Identifies the SLA Profile for this subscriber host (optional).

These strings are used for a lookup in one or more maps to find the names of the sub-profile and sla-profile to use. If none of the maps contained an entry for these strings, the names will be determined based on a set of defaults.

Only when the names for both the sub-profile and sla-profile are known, the subscriber host can be instantiated. If even no default is found for either profile, the DHCP ACK is dropped and the host will not gain network access.

9.3.7.3 Determining the Subscriber Profile

All hosts (devices) belonging to the same subscriber will be subject to the same HQoS processing. The HQoS processing is defined in the sub-profile. A sub-profile refers to an existing scheduler policy and offers the possibility to overrule the rate of individual schedulers within this policy.

Because all subscriber hosts of one subscriber use the same scheduler policy instance, they must all reside on the same I/O module.

The figure below shows how the sub-profile is derived, based on the sub-ident string, the sub-profile string and/or the provisioned data structures. The numbers associated with the arrows pointing toward the subscriber profiles indicate the precedence of the checks.

Figure 76 **Determining the Subscriber Profile**

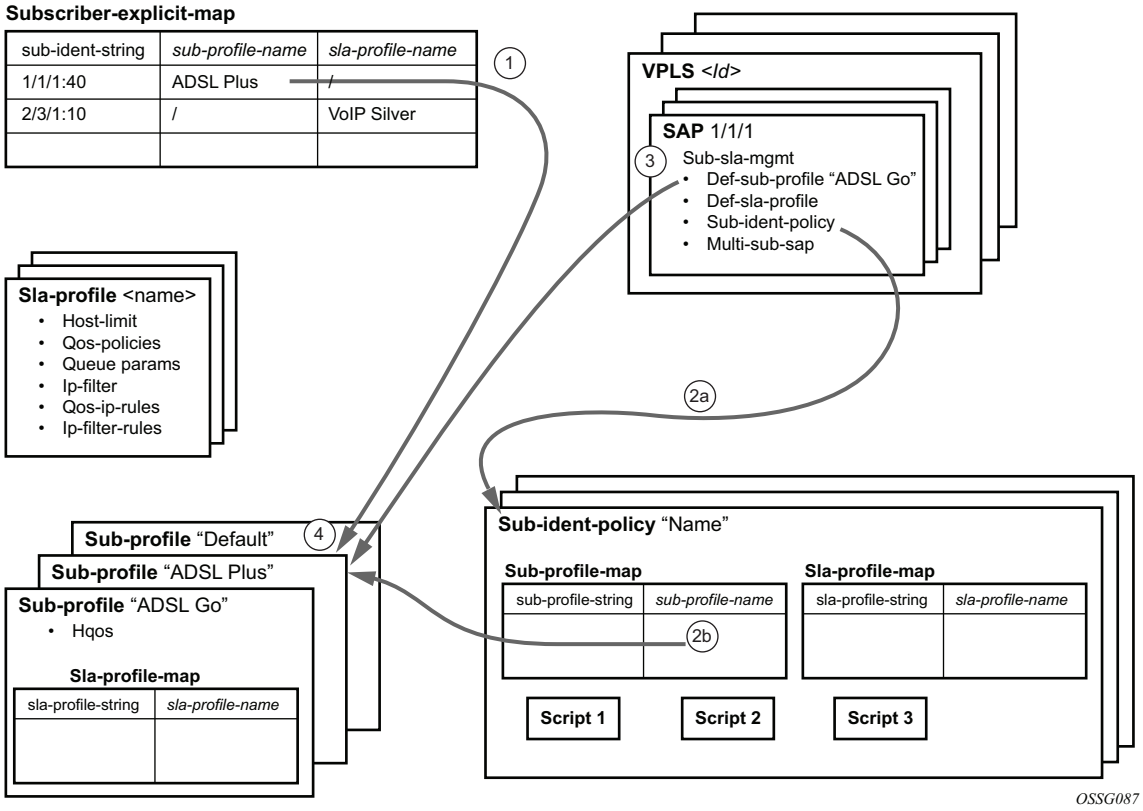

1. A lookup in the **explicit-subscriber-map** is done with the sub-ident string returned by the script. If a matching entry is found, the sub-profile-name (if defined) is taken. Otherwise:

2. If a **sub-ident-policy** is defined on the SAP, a lookup is done on its **sub-profile-map** with the sub-profile string from the script. The sub-profile-name is taken from the entry.

   If no entry was found, then:

3. If provisioned, the sub-profile-name is taken from the **def-sub-profile** attribute on the SAP. If not provisioned, then:
4. The **sub-profile** with the name “default” is selected (if provisioned). If this is not provisioned, there are no other alternatives, the ACK is dropped, and the host will not gain access.

9.3.7.4 Determining the SLA Profile

For each host that comes on-line, the router also needs to determine which SLA profile to use. The SLA profile will determine for this host:

- The QoS-policies to use:
  - classification
  - queues/policers
  - queue mapping
- The egress scheduling policies to use:
  - egress HQoS
- The IP filter to use.

The SLA profile also has a **host-limits** attribute which limits the number of hosts (belonging to the same subscriber) on a certain SAP that can be using this SLA profile.

The classification and the queue mapping are shared by all the hosts on the same forwarding complex that use the same QoS policy (by their SLA profile).

The queues/policers are shared by all the hosts (of the same subscriber) on the same SAP that are using the same SLA profile. In other words, queues/policers are instantiated when, on a given SAP, a host of a subscriber is the first to use a certain SLA profile. This instantiation is referred to as an SLA profile instance. Ingress queues can be parented to a scheduler referenced in the ingress of a subscriber profile. Egress policers and queues can be parented to a scheduler referenced in the egress of a subscriber or SLA profile, or to a port scheduler.

A scheduler policy can be applied to the egress an SLA profile, allowing its schedulers to be the parent for its queues and for its tier 1 schedulers to be parented to a scheduler in a scheduler policy applied to the egress of a subscriber profile or a Vport, or to a port scheduler applied to a port or Vport. Configuring scheduler overrides is allowed for SLA profile egress schedulers. The configuration of a scheduler policy in the egress of an SLA profile is supported for all host types only on Ethernet interfaces. It is not supported for ESM over MPLS pseudowires, nor is HQoS adjustment and host tracking supported on its schedulers.

The following show, monitor and clear commands are available related to the SLA profile scheduler:

```
show qos scheduler-hierarchy subscriber sub-ident-string sla-profile sla-profile-
name
sap sap-id [scheduler scheduler-name] [detail]
```

The **show qos scheduler-hierarchy subscriber** command (shown above) displays the scheduler hierarchy with the SLA profile scheduler as the root. Note that if the SLA profile scheduler is orphaned (that is when the scheduler has a parent which does not exist) then the hierarchy is only shown when the show command includes the **sla-profile** and **sap** parameters.

If the SLA profile scheduler is orphaned (that is when the scheduler has a parent which does not exist) then the hierarchy is only shown when the show command includes the sla-profile and SAP parameters.

```
monitor qos scheduler-stats subscriber sub-ident-string [interval seconds] [repeat
repeat] [absolute|rate] sap sap-id sla-profile sla-profile-name

show qos scheduler-stats subscriber sub-ident-string sap sap-id sla-profile sla-
profile-name [scheduler scheduler-name]

clear qos scheduler-stats subscriber sub-ident-string sap sap-id sla-
profile sla-profile-name [scheduler scheduler-name]
```

Figure 77 shows a graphical description of how the SLA profile is derived based on the subscriber identification string, the SLA profile string and the provisioned data structures. The numbers on the arrows towards the SLA profile indicate the *priority* of the provisioning (the lower number means the higher priority).

Figure 77 **Determining the SLA Profile**

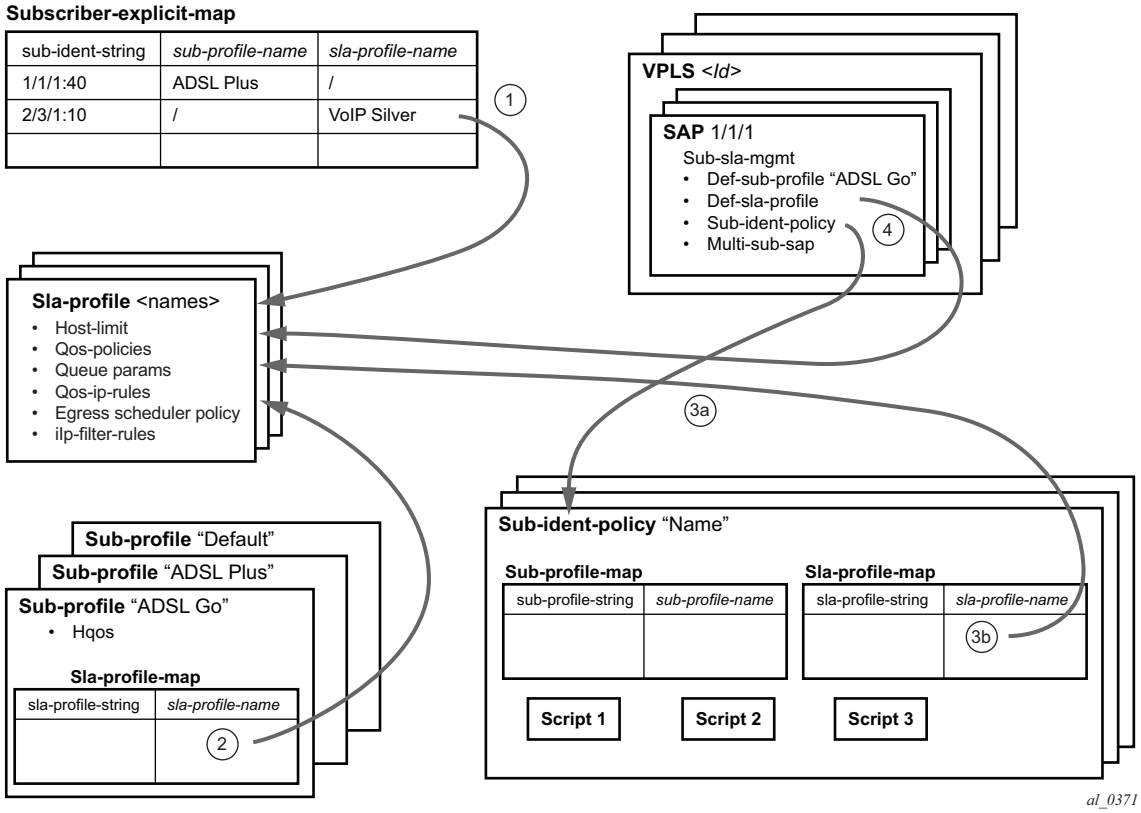

1. A lookup is done with the sub-ident string returned by the script in the **explicit-subscriber-map**. If a matching entry is found, the sla-profile-name is taken from it – if defined. Otherwise:
2. A lookup with the sla-profile string from the script is done in the **sla-profile-map** of the sub-profile found earlier. The sla-profile-name from the found entry is taken. If no entry was found, then:
3. A lookup is done with the sla-profile string in the **sla-profile-map** of the **sub-ident-policy** configured on the SAP. The sla-profile-name from the found entry is taken. If no **sub-ident-policy** was configured on the SAP or no entry was found, then:
4. If provisioned, the sla-profile-name is taken from the **def-sla-profile** attribute on the SAP. If not provisioned, there are no more alternatives, the ACK is dropped, and the host will not gain access.

9.3.7.4.1 SLA Profile Instance Sharing

Each subscriber host or session has an SLA Profile Instance (SPI) associated with it. The SPI, is by default, determined by the subscriber ID, the SLA profile name, and the SAP where the subscriber host or session is active. See Figure 72.

SPIs with the same SLA profile name, have the same configuration, however, the following functions are effective per SPI:

- enforcing the different host limits
- instantiation of queues and policers
- accounting statistics
- credit control functions

For a bridged Residential Gateway deployment, typically multiple IPoE or PPPoE sessions per subscriber are active on the BNG. The next sections describe the different SPI sharing mechanisms that apply for multiple subscriber sessions from the same subscriber, that are active on the same SAP with the same SLA profile name assigned.

SPI Sharing per SAP

By default, all subscriber sessions or hosts from the same subscriber, active on the same SAP and with the same SLA profile assigned, share an SPI. The default SPI sharing is per SAP, as depicted in Figure 78.

Figure 78 **SLA Profile Instance per SAP**

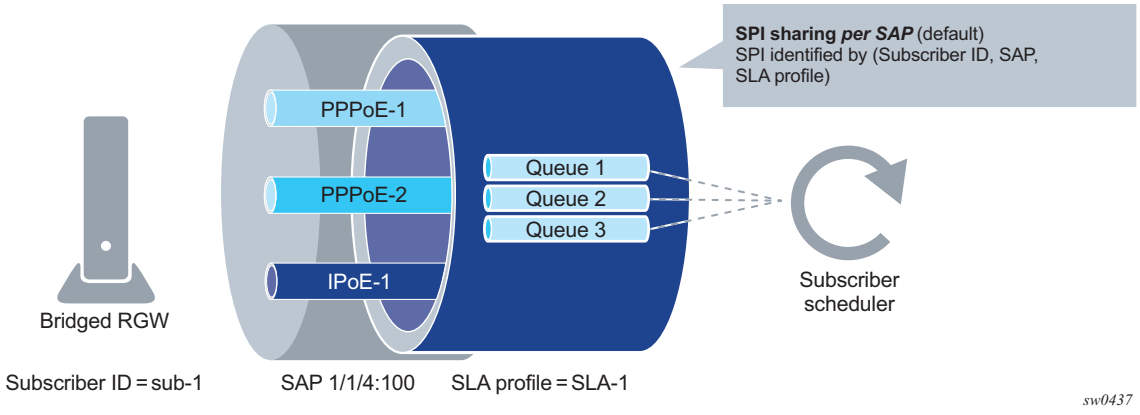

With SPI sharing per SAP, traffic from all subscriber sessions on a given SAP and with the same SLA profile associated are mapped to the same set of queues and policers for QoS handling. Statistics from these queues and policers are also used in accounting. Per-host or per-session accounting modes cannot report counters for individual sessions unless their traffic is mapped in separate queues.

SPI sharing per SAP is the default configuration in an SLA profile and applies to PPPoE sessions, IPoE sessions (enabled on the group-interface) and IPoE hosts (IPoE sessions are disabled on the group-interface):

```
configure
    subscriber-mgmt
        sla-profile "sla-profile-1"
            def-instance-sharing per-sap
```

SPI Sharing - per Session

If QoS handling or accounting per-IPoE or per-PPPoE session is required, then the SPI sharing is configured to per-session sharing in the SLA profile:

```
configure
    subscriber-mgmt
        sla-profile "sla-profile-1"
            def-instance-sharing per-session
```

Per-session sharing applies to PPPoE sessions and IPoE sessions (enabled on the group interface). An IPoE host setup fails when IPoE sessions are disabled on the group interface and per-session sharing is configured.

Each IPoE or PPPoE session from the same subscriber, active on the same SAP and having the same SLA profile assigned, has its own set of queues and policers. Per-session SPI sharing is depicted in Figure 79.

Figure 79 **SLA Profile Instance per Session**

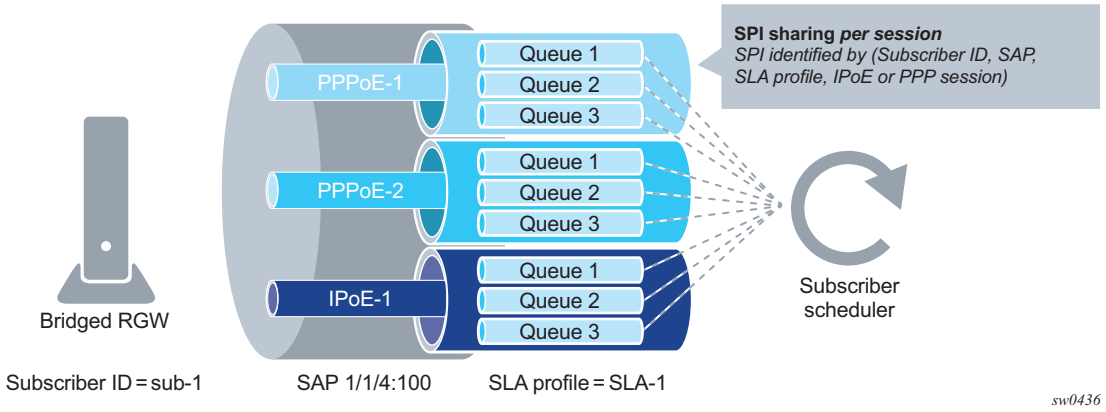

Note: SPI sharing per session is not supported on HS MDA and on HSQ with **hs-sla-mode single.**

SPI per Group

When even more granular control is needed over which sessions share an SPI, an SPI sharing group identifier can be specified during IPoE or PPPoE session authentication. This overrides the default SPI sharing method for that session as configured in the SLA profile.

Per-group SPI sharing is depicted in Figure 80. The same SPI is shared by all IPoE and PPPoE sessions from the same subscriber, active on the same SAP, having the same SLA Profile assigned and having the same SPI sharing group identifier.

Note: SPI sharing per group is not supported on HS-MDA and on HSQ with **hs-sla-mode single.**

Figure 80 **SLA Profile Instance per Group**

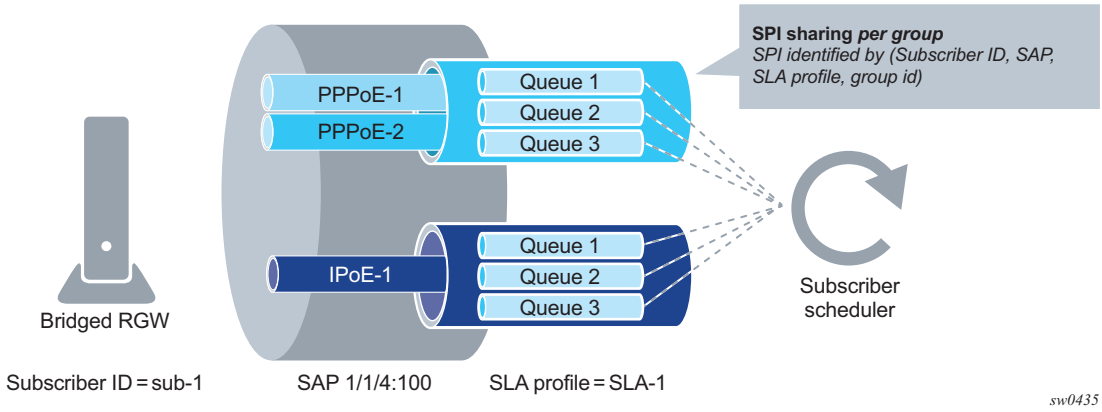

The SPI sharing group identifier is an integer value in the range 0 to 65535 and can be specified in authentication using:

- A local user database lookup:

```
configure
   subscriber-mgmt
      local-user-db local-user-db-name
         ipoe | pppoe
            host host-name
               identification-strings
                  spi-sharing-group-id <group-id>
```

  Configure **no spi-sharing-group-id** to apply the **def-instance-sharing** method as configured in the SLA profile.

- RADIUS, by including the [241.26.6527.47] Alc-SPI-Sharing-Id VSA in an Access-Accept message:
  - Value "group:<group-id>" to enable SPI sharing per group identifier
  - Value "default" to apply the **def-instance-sharing** method as configured in the SLA profile

  Refer to the *7750 SR and VSR RADIUS Attributes Reference Guide* for a detailed description of the attribute.

- Diameter NASREQ, by including the Vendor specific [NOKIA-1036] Alc-SPI-Sharing grouped AVP in an AA-Answer message:

```
Alc-SPI-Sharing ::= < AVP Header: 1036 >
                        {Alc-SPI-Sharing-Type}
                        [Alc-SPI-Sharing-Id]
```

  - To enable SPI sharing per group identifier
    - [NOKIA-1037] Alc-SPI-Sharing-Type = 2
    - [NOKIA-1038] Alc-SPI-Sharing-Id = <group-id>

– To apply the **def-instance-sharing** method as configured in the SLA profile:

• [NOKIA-1037] Alc-SPI-Sharing-Type = 0

See the Diameter and Diameter Applications, AA-Answer Message — Accepted Authorization AVPs section for a detailed description of the attribute.

• Diameter Gx, by including the Vendor specific [NOKIA-1036] Alc-SPI-Sharing grouped AVP in a CCA message:

```
Alc-SPI-Sharing ::= < AVP Header: 1036 >
                        {Alc-SPI-Sharing-Type}
                        [Alc-SPI-Sharing-Id]
```

– To enable SPI sharing per group identifier

• NOKIA-1037] Alc-SPI-Sharing-Type = 2

• NOKIA-1038] Alc-SPI-Sharing-Id = <group-id>

– To apply the **def-instance-sharing** method as configured in the SLA profile:

• [NOKIA-1037] Alc-SPI-Sharing-Type = 0

See *7750 SR Gx AVPs Reference Guide* for a detailed description of the attribute.

• Python:

– alc.dts.setESM module: alc.dtc.SpiSharingGroupId = <group-id>

– alc.esm.set module: alc.esm.SpiSharingGroupId = <group-id>

See the DHCP Management, ESM-Related Python Variables section for further details.

Per-group sharing applies to PPPoE sessions and IPoE sessions (enabled on the group interface). An IPoE host setup fails when IPoE sessions are disabled on the group interface and an SPI sharing group identifier is specified.

Dynamic Changes of SLA Profile and SPI Sharing

During the lifetime of an IPoE or PPPoE session, the SLA profile and the SPI sharing can change. Such a dynamic change can be triggered by re-authentication, RADIUS CoA, or Diameter Gx RAR by specifying a new SLA profile and optionally, an SPI sharing group ID.

Table 76 describes the different transitions in SPI sharing because of re-authentication, RADIUS CoA, or Diameter Gx RAR.

Table 76 **Dynamic Changes of SPI Sharing**

from ↓ **to**	**SLA Profile and SPI sharing info provided for dynamic change**
per sap ↓ per sap	SLA profile = <SLA profile name> • SLA profile with "def-instance-sharing per-sap" [SPI sharing type = default] • Optional. Value must be default when present • SPI sharing ID must not be present
per session ↓ per session	[SLA profile = <SLA profile name>] • SLA profile with "def-instance-sharing per-session" [SPI sharing type = default] • Optional. Value must be default when present • SPI sharing ID must not be present
per group ↓ per group	SLA profile = <SLA profile name> • Optional. Current SLA profile name is used when not specified SPI sharing type = group SPI sharing ID = <group-id> • Overrides the "def-instance-sharing" configured in the SLA profile
per sap ↓ per group	SLA profile = <SLA profile name> • Optional. Current SLA profile name is used when not specified SPI sharing type = group SPI sharing ID = <group-id> • Overrides the "def-instance-sharing" configured in the SLA profile
per session ↓ per group	SLA profile = <SLA profile name> • Optional. Current SLA profile name is used when not specified SPI sharing type = group SPI sharing ID = <group-id> • Overrides the "def-instance-sharing" configured in the SLA profile

Table 76 **Dynamic Changes of SPI Sharing (Continued)**

from ↓ **to**	**SLA Profile and SPI sharing info provided for dynamic change**
per group ↓ per sap	SLA profile = <SLA profile name> • Optional. Current SLA profile name is used when not specified • SLA profile with "def-instance-sharing per-sap" SPI sharing type = default • SPI sharing ID must not be present --- SLA profile = <SLA profile name> • SLA profile with "def-instance-sharing per-sap"
per group ↓ per session	[SLA profile = <SLA profile name>] • Optional. Current SLA profile name is used when not specified • SLA profile with "def-instance-sharing per-session" SPI sharing type = default • SPI sharing ID must not be present --- SLA profile = <SLA profile name> • SLA profile with "def-instance-sharing per-session"
per sap ↓ per session	SLA profile = <SLA profile name> • SLA profile with "def-instance-sharing per-session" [SPI sharing type = default] • Optional. Value must be default when present • SPI sharing ID must not be present
per session ↓ per sap	SLA profile = <SLA profile name> • SLA profile with "def-instance-sharing per-sap" [SPI sharing type = default] • Optional. Value must be default when present • SPI sharing ID must not be present

Identifying the SPI

An SPI is uniquely identified by the following characteristics:

- the subscriber identifier
- the SAP on which the subscriber session is active
- the SLA profile name
- An SPI sharing identifier that has two parts to support overlapping ids between groups and sessions:
  - SPI sharing type: per SAP, per IPoE session, per PPP session, per group
  - SPI sharing id:
    - an integer value determined by the system for SPI sharing per session
    - an integer value in the range 0 to 65535 specified by the user for SPI sharing per group
    - not required for SPI sharing per SAP

Below are some examples for SPI representations in the system:

- SPI sharing per SAP

```
A:PE-1# show service active-subscribers detail
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber sub-01 (sub-profile-12)
-------------------------------------------------------------------------------
--- snip---
-------------------------------------------------------------------------------
(1) SLA Profile Instance
    - sap:[1/1/4:1201.41] (IES 1000 - group-int-1-1)
    - sla:sla-profile-12
-------------------------------------------------------------------------------
--- snip---
A:PE-1# show service active-subscribers hierarchy
===============================================================================
Active Subscribers Hierarchy
===============================================================================
-- sub-01 (sub-profile-12)
   |
   +-- sap:[1/1/4:1201.41] - sla:sla-profile-12
       |
       |-- PPP-session - mac:00:51:00:00:01:41 - sid:1 - svc:1000
       |   |
       |   +-- 10.1.1.141 - IPCP
       |
       |-- PPP-session - mac:00:51:00:00:01:42 - sid:2 - svc:1000
       |   |
       |   +-- 10.1.1.142 - IPCP
       |
       +-- PPP-session - mac:00:51:00:00:01:43 - sid:3 - svc:1000
           |
```

```
          +-- 10.1.1.143 - IPCP
-------------------------------------------------------------------------------
Number of active subscribers : 1
Flags: (N) = the host or the managed route is in non-forwarding state
===============================================================================
*A:PE-1# show qos scheduler-hierarchy subscriber "sub-01" detail
===============================================================================
Scheduler Hierarchy - Subscriber sub-01
===============================================================================
--- snip ---
Root (Egr)
| slot(1)
|--(Q) : Sub=sub-01:sla-profile-12 1000->1/1/4:1201.41->6  (Port 1/1/4)
|   |    AdminPIR:1500       AdminCIR:1500
--- snip ---
```

- SPI sharing per Session

Note: Although they can have the same value as in the following output, an SPI sharing id is not the same as the PPP session id.

```
A:PE-1# show service active-subscribers detail
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber sub-01 (sub-profile-12)
-------------------------------------------------------------------------------
---snip---
-------------------------------------------------------------------------------
(1) SLA Profile Instance
    - sap:[1/1/4:1201.41] (IES 1000 - group-int-1-1)
    - sla:sla-profile-12 PPP session:10
-------------------------------------------------------------------------------
---snip---
A:PE-1# show service active-subscribers hierarchy
===============================================================================
Active Subscribers Hierarchy
===============================================================================
-- sub-01 (sub-profile-12)
   |
   |-- sap:[1/1/4:1201.41] - sla:sla-profile-12 PPP session:10
   |   |
   |   |-- PPP-session - mac:00:51:00:00:01:41 - sid:10 - svc:1000
   |   |   |
   |   |   +-- 10.1.1.141 - IPCP
   |
   |-- sap:[1/1/4:1201.41] - sla:sla-profile-12 PPP session:11
   |   |
   |   |-- PPP-session - mac:00:51:00:00:01:42 - sid:11 - svc:1000
   |   |   |
   |   |   +-- 10.1.1.142 - IPCP
   |
   +-- sap:[1/1/4:1201.41] - sla:sla-profile-12 PPP session:12
       |
```

```
       +-- PPP-session - mac:00:51:00:00:01:43 - sid:12 - svc:1000
           |
           +-- 10.1.1.143 - IPCP
-------------------------------------------------------------------------------
Number of active subscribers : 1
Flags: (N) = the host or the managed route is in non-forwarding state
===============================================================================
A:PE-1# show qos scheduler-hierarchy subscriber "sub-01" detail
===============================================================================
Scheduler Hierarchy - Subscriber sub-01
===============================================================================
--- snip ---
Root (Egr)
| slot(1)
|--(Q) : Sub=sub-01:sla-profile-12:PPP-10 1000->1/1/4:1201.41->6  (Port 1/1/4)
|   |    AdminPIR:1500       AdminCIR:1500
--- snip ---
```

- SPI sharing per-group

```
A:PE-1# show service active-subscribers detail
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber sub-01 (sub-profile-12)
-------------------------------------------------------------------------------
---snip---
-------------------------------------------------------------------------------
(1) SLA Profile Instance
    - sap:[1/1/4:1201.41] (IES 1000 - group-int-1-1)
    - sla:sla-profile-12 group:100
-------------------------------------------------------------------------------
---snip---
A:PE-1# show service active-subscribers hierarchy
===============================================================================
Active Subscribers Hierarchy
===============================================================================
-- sub-01 (sub-profile-12)
   |
   |-- sap:[1/1/4:1201.41] - sla:sla-profile-12 group:100
   |   |
   |   |-- PPP-session - mac:00:51:00:00:01:41 - sid:15 - svc:1000
   |   |   |
   |   |   +-- 10.1.1.141 - IPCP
   |   |
   |   |-- PPP-session - mac:00:51:00:00:01:42 - sid:14 - svc:1000
   |   |   |
   |   |   +-- 10.1.1.142 - IPCP
   |
   +-- sap:[1/1/4:1201.41] - sla:sla-profile-12 group:200
       |
       +-- PPP-session - mac:00:51:00:00:01:43 - sid:13 - svc:1000
           |
           +-- 10.1.1.143 - IPCP
-------------------------------------------------------------------------------
Number of active subscribers : 1
Flags: (N) = the host or the managed route is in non-forwarding state
===============================================================================
```

```
A:PE-1#  show qos scheduler-hierarchy subscriber "sub-01" detail
===============================================================================
Scheduler Hierarchy - Subscriber sub-01
===============================================================================
---snip---
Root (Egr)
| slot(1)
|--(Q) : Sub=sub-01:sla-profile-12:Group-100 1000->1/1/4:1201.41->6  (Port 1/1/4)
|   |    AdminPIR:1500       AdminCIR:1500
---snip---
```

In RADIUS accounting messages, the SPI is uniquely defined by the following attributes:

```
configure
    subscriber-mgmt
        radius-accounting-policy "acct-policy-1"
            include-radius-attribute
                subscriber-id
                nas-port-id
                sla-profile
                spi-sharing
```

For example:

```
 NAS PORT ID [87] 13 1/1/4:1201.41
    VSA [26] 40 Nokia(6527)
      SUBSC ID STR [11] 6 sub-01
      SLA PROF STR [13] 14 sla-profile-12

    # SPI sharing per SAP:
    VSA [241.26] 3 NOKIA(6527)
      SPI SHARING_ID [47] 3 SAP
    # SPI sharing per session:
    VSA [241.26] 14 NOKIA(6527)
      SPI SHARING_ID [47] 14 PPP session:12
    # SPI sharing per group:
    VSA [241.26] 14 NOKIA(6527)
      SPI SHARING_ID [47] 9 group:100
```

9.3.7.4.2 SLA-Based Egress QoS Marking

The egress QoS marking for subscriber-host traffic is derived from the SAP-egress QoS policy associated with a corresponding SAP, rather than from the SLA profile associated with the corresponding subscriber host. As a consequence, no egress QoS marking (Dot1p marking is set to 0, the dscp/prec field is kept unchanged) is performed for traffic transmitted on a managed SAP because by default, sap-egress policy 1 is attached to every managed SAP.

The default value of the “qos-marking-from-sap” flag is enabled. This means that the qos-marking defined in the SAP egress QoS policy associated with the SAP will be used. The default setting of this flag in a combination with managed-SAP will result in the same behavior as in the current system (dot1p=0, dscp/prec is unchanged).

If the **no qos-marking-from-sap** command is executed, then both the Dot1p marking and DSCP marking are derived from the sla-profile.

Changing the flag setting in the SLA profile being used by any subscriber-hosts (this includes subscriber-hosts on managed-SAPs as well) will be allowed.

The following MC traffic characteristics apply:

- On Layer 3 subscriber interfaces, MC is not supported so it is impossible to enable it at the SAP level or at the sla-instance level.
- On Layer 2 SAPs IGMP snooping is supported while it is not supported on the sla-instance level. Therefore, any MC traffic transmitted at egress belongs to a SAP (meaning it will use SAP queues), rather than to a sla-instance.
- The special case are SAPs with a profiled-traffic-only flag enabled. Although it is possible to define an sla-profile applicable to a Layer 2-host, this will not be taken as reference for marking mc-traffic, but rather SAP settings will be used.

9.3.8 Auto-Sub ID

The subscriber ID name (*sub-id*) is a mandatory object that binds all hosts of a given subscriber together. Briefly, the *sub-id* name represents a residential household. Many management/troubleshooting and even billing operations rely on the *sub-id* name entity. The *sub-id* name is required for the host creation process, and it can be supplied by any authentication source, such as RADIUS, Diameter, LUDB, or Python. It can be derived from the *sap-id* or can be statically provisioned in the form of a string.

In some ESM deployments, it is desirable that the *sub-id* is automatically generated within the router rather than burdening the OSS with this function. A typical application for auto sub-id is as follows:

- RADIUS server provides the SLA profile string and the sub-profile string but not the sub-id string.
- The *sub-id* name is automatically generated and formatted based on the configured options.

The following are the properties of auto sub-id generation:

- The automatic generation of the subscriber ID name can be based on any combination of the following fields:
  - MAC address
  - sap-id
  - circuit-id
  - remote-id
  - session-id

There can be only a single set of subscriber identification fields defined per host type (IPoE or PPPoE) per chassis. If the combination of the fields must be modified, the existing subscribers with an automatically generated subscriber ID must be manually terminated. Considering that remote termination of the IPoE subscribers by a DHCP server is not supported by all DHCP client vendors through the FORCERENEW DHCP message (RFC 3203, *DHCP reconfigure extension*), changing the subscriber fields while subscribers with automatically generated subscriber ID are active should be avoided.

The subscriber ID name automatic generation takes place at the end of the host initiation process (after the authentication phase is completed) and only in case whereby the subscriber ID had not been already provided by any other more specific means (RADIUS, Diameter, LUDB, or Python).

The format of the sub-id name can be either a 10-character encoded string (characters A to Z and 0 to 9) or a user- friendly string based on the subscriber identification fields. The maximum length of the subscriber ID name is 32 characters.

The subscriber ID name will not be passed in the Access-Request to the RADIUS server since it is generated after the authentication phase.

The subscriber ID name can be automatically generated regardless of how the SLA or subscriber profile strings are obtained (RADIUS, LUDB, Python, or static).

The subscriber identification fields used in automatic generation of the subscriber ID name are enabled at the system level.

```
configure
subscriber-mgmt
     auto-sub-id-key
          ppp-sub-id-key [mac] [sap-id] [circuit-id]
             [remote-id] [session-id]
          ipoe-sub-id-key [mac] [sap-id] [circuit-id]
             [remote-id]
```

If no *sub-id-key* per host type is configured, the defaults are:

PPPoE host type:	mac, sap-id, session-id
IPoE host type:	<mac, sap-id>.

The order in which the fields are configured is important because the subscriber ID name potentially becomes a concatenated string of the subscriber host identifiers in the order in which they are provisioned. The subscriber ID cannot be longer than 32 characters.

- If the length of the concatenated fields for the subscriber ID name is longer than 32 characters, the host creation will fail.
- If the circuit ID or remote ID is in the key and they contain non-printable characters, their place in subscriber ID name will be formatted in hex instead of ASCII. ASCII printable characters can contain the byte values 0x20 to 0x7E. All other values are ASCII non-printable and thus, are formatted in hex characters.

The following would generate a subscriber ID name: xx:xx:xx:xx:xx:xx|1/1/3:23|44. The length of such subscriber ID name would be 29B.

- mac: xx:xx:xx:xx:xx:xx
- sap: 1/1/3:23
- session-id: 44 (16bits length)

If the key contains the circuit ID as: 0x610163 (3 bytes), then the subscriber ID name will be formatted as 610161, in hex, since 01 hex is non-printable in ASCII. Then the subscriber ID name's length is 6B.

However, if the circuit ID is 0x616263 (3 bytes), then the string will be formatted as ASCII string abc (three characters). The subscriber ID name's length is 3B.

The assignment of the subscriber ID to dynamic hosts is performed in the following order:

1. From authentication sources: RADIUS, Diameter, LUDB, or Python.

   A subscriber ID name obtained from authentication sources can conflict with the format of an implicit auto-generated subscriber ID name. When this happens, the subscriber host or session setup fails. Therefore, when implicit subscriber ID name generation is enabled (the default), a 10-character string containing the characters A to Z and 0 to 9 should not be returned from authentication sources. Information in Step 3 describes information to disable the implicit automatic generation of a subscriber ID name.

2. An explicit configured default is configured as *def-sub-id*:
   - At the SAP level for static SAPs:

```
configure
   service ies/vprn
      subscriber-interface <sub-if-name>
         group-interface <grp-if-name>
            sap <sap-id>
               sub-sla-mgmt
                  def-sub-id use-sap-id | use-auto-id | string <sub-id>
```

   - In the MSAP policy for managed SAPs:

```
configure
   subscriber-mgmt
      msap-policy <name>
         sub-sla-mgmt
            def-sub-id use-sap-id | use-auto-id | string <sub-id>
```

   where:
   - **use-sap-id**: the *sub-id* name is the SAP identifier
   - **use-auto-id**: the *sub-id* name is a combination of the identifiers specified in **auto-sub-id-key**.

     The *sub-id* name is in a readable format, that is, a concatenation of the fields in the **pppoe-sub-id-key** or **ipoe-sub-id-key** command separated by a "|" character.
   - **string**: the *sub-id* name is a user defined string

3. Implicitly generated default:

When no subscriber ID name is provided in authentication, and no explicit default is configured, then the system, by default, automatically generates a subscriber ID name, a 10-character string, using characters A to Z and 0 to 9, that is based on the fields defined in the:

- **ppp-sub-id-key** command for PPP host types. If no such fields are explicitly defined, the default are assumed: mac, sap-id, session-id.
- **ipoe-sub-id-key** command for IPoE host types. If no such fields are explicitly defined, the defaults are assumed: mac, sap-id .

The implicitly generated subscriber ID name is unique per chassis as well as in dual-homed environments.

With implicit automatic subscriber ID name generation enabled, neither the *old-sub-ident-string* value nor the *new-sub-ident-string* value in the **tools>perform>subscriber-mgmt>re-ident-sub** *old-sub-ident-string* **to** *new-sub-ident-string* CLI command can refer to a subscriber ID in an implicitly generated default format.

The implicit automatic subscriber ID name generation can be disabled with the following command: **configure>subscriber-mgmt>auto-sub-id-key>no implicit-generation**.

- The implicit *auto-sub-id* name generation cannot be disabled when there are active subscribers in the system with an implicit automatically generated subscriber ID name.
- When disabled, the implicit *auto-sub-id* name generation cannot be enabled when there are active subscribers in the system.

With implicit subscriber ID generation disabled, the subscriber host or session setup fails when no subscriber ID name is provided in authentication, and no explicit default is configured. A 10-character subscriber ID name format, using the characters A to Z and 0 to 9, can be returned from authentication sources without risk of conflict.

Static subscribers are required to have the sub-id manually configured.

9.3.8.1 Sub-id Identifiers

The sub-id can be based on any combination of the following identifiers:

- The sap-id, in combination with any other allowable identifier, will be used as the search key. This assumes a 1:1 (subscriber per SAP) deployment model.

- The circuit-id, in combination with any other allowable identifier, will be used to identify subscribers. This can be used in 1:1 deployment model, or in service per SAP deployment model. Circuit-id is applicable to IPoE v4 type hosts (option 82), to IPoE v6 type hosts (option 18 – interface-id) and PPPoE hosts (remote agent option signaled by PPPoE tags). The format of circuit-id is identical for IPv4 and IPv6 hosts.
- The remote-id, in combination with any other allowable identifier, will be used to identify subscribers. This can be used in 1:1 deployment model, or in service per SAP deployment model. The remote-id is applicable to IPoE v4 type hosts (option 82), to IPoE v6 type hosts (option 37) and PPPoE hosts (remote agent option signaled via PPPoE tags).
- The mac address (in combination with any other allowable identifier will be used to identify subscribers. This assumes a 1:1 deployment model.
- The PPPoE session id, in combination with any other allowable identifier, is applicable only to PPPoE hosts. The session-id used will be of the first host that is instantiated for the subscriber.

9.3.8.2 Dual Stack Hosts

Auto-generation of sub-id names for subscribers with a single dual stack hosts (IPoE and PPPoE) is enabled by default by not explicitly provisioning anything for the def-sub-id. The sub-id name would be semi-randomly generated based on the <mac, sap-id, session-id> for PPPoE hosts and the <mac, sap-id> combination for IPoE host.

9.3.8.3 Mixing Hosts with Auto-Generated IDs and non Auto-Generated IDs

Hosts with different sub-id names but identical auto-sub-id keys are not linked into the same subscriber. Such scenarios can arise with hosts with the same auto-sub-id keys but different methods for obtaining the sub-id name. For example, one host relying on auto-generated sub-id name while the other is using explicit configuration methods (sap-id, string, RADIUS or LUDB). If the auto-generated sub-id name and explicit sub-id name are the same, the host will be tied into the same subscriber.

For example:

The default auto-sub-id for the following two hosts are <mac, sap-id>.

Host X on SAP 1/1/1:1 with MAC 00:00:00:00:00:01 obtains sub-id through RADIUS.

Host Y on SAP 1/1/1:1 with MAC 00:00:00:00:00:01 has sub-id auto-generated.

Regardless of which host comes up first, those two hosts at the end will belong two different subscribers as long as their sub-ids are different.

9.3.8.4 PPPoA/PPPoEoA Considerations

PPPoA/PPPoEoA hosts will adhere to the same rules as PPPoE. Fields that are supported in PPPoE but not PPPoA/PPPoEoA will be simply ignored.

Fields that are not supported in PPPoA are:

- Remote-id
- Circuit-id
- MAC address

Fields that are most likely not applicable to PPPoEoA are:

- Remote-id
- Circuit-id

9.3.8.5 Deployment Considerations

The following is a possible deployment example scenario.

```
config
 subscriber-mgmt
     auto-sub-id-key
          ppp-sub-id-key sap-id
          ipoe-sub-id-key mac circuit-id



config
 service vprn 10
      subscriber-interface <sub-if-name>
      authentication-policy <auth-pol-name>
      group-interface <grp-if-name>
          sap 1
               sub-sla-mgmt
                    def-sub-id use-sap-id
                    sub-ident-policy <ident-pol-name>
          sap 2
```

```
        sub-sla-mgmt
             def-sub-id auto-id
             sub-ident-policy <ident-pol-name>
   sap 3
        sub-sla-mgmt
             def-sub-id "sub3"
             sub-ident-policy <ident-pol-name>
   sap 4
        sub-sla-mgmt
                  sub-ident-policy <ident-pol-
                    name>
```

Assume the following cases:

1. RADIUS returns the sub-id on all four SAPs.
2. RADIUS does not return the sub-id string on any of the SAPs.

In the first case where RADIUS returns the sub-id string, the following will occur:

- On all 4 SAPs the sub-id string will be assigned by the RADIUS server. Defaults have no effect, and neither do identifiers specified under the auto-sub-id-key node.

In the second case, the effects are the following:

- On SAP1 the sub-id name will be the <sap-id> (1/1/1:3)
- On SAP 2 the sub-id name will be <sap-id> for PPPoE hosts and <mac>-<circuit-id> concatenation for IPoE type hosts.

For example:

1/1/1:100 for PPPoE

AC:AB:AA:AD:AE:AE-AN-id eth 1/1/1/1:2 for IPoE

(circuit-ID format is: Access-Node-Identifier atm slot/port:vpi.vci or Access-Node-Identifier eth slot/port:[vlan-id]).

Note that the circuit-ID can itself be 63B in length whereas the length of the sub-id name is limited to 32 Bytes. In the case noted above, the sub-id name length would be 38 Bytes (>32B) and the host instantiation would fail.

- On SAP3 the sub-id name will be the literal ‘sub3’ for PPPoE and IPoE hosts.
- On SAP4 the sub-id name will be a semi-random value based on <sap-id> for PPPoE hosts and the <mac, circuit-id> combination for IPoE hosts.

9.3.8.6 Caveats

Only a single combination of the subscriber fields used to auto generate sub-id is allowed per host type (IPoE or PPPoE) and per chassis. In case that the combination of the fields needs to be changed, the existing subscribers with an auto-generated sub-id must be manually terminated. Considering that remote termination of the IPoE subscribers by DHCP server is not supported by all DHCP client vendors through FORCERENEW DHCP message (RFC 3203), changing the subscriber fields while subscribers with auto generated sub-id are active should be avoided.

9.3.9 Limiting Subscribers and Hosts on a SAP

A number of configuration parameters are available to control the maximum amount of subscribers and/or hosts that can be simultaneously active on a SAP:

- multi-sub-sap — Limits the number of subscribers (dynamic + static) on a SAP
- lease-populate — Limits the number of dynamic hosts on a SAP
- host-limits — Limits the number of hosts (dynamic + static) per SPI.

If any of these limits are reached, a new host will be denied access. The only exception is when the **host-limits** command is configured with the keyword **remove-oldest** specified, then the oldest active host is dropped and the new host is granted access. The dynamic host with the least remaining lease time will be considered the oldest host. Note that **remove-oldest** only applies to DHCP and ARP-host, and it is not for PPPoE or IPv6 ESM hosts.

9.3.10 Static Subscriber Hosts

While it is typically preferred to have all hosts provisioned dynamically through DHCP snooping, it may be needed to provide static access for specific hosts (those that do not support DHCP).

Since a subscriber identification policy is not applicable to static subscriber hosts, the subscriber identification string, subscriber profile and SLA profile must be explicitly defined with the host's IP address and MAC address (if Enhanced Subscriber Management is enabled).

If an SPI associated with the named SLA profile already exists on the SAP for the subscriber, the static subscriber host is placed into that SPI. If an SPI does not yet exist, one will be created if possible. If the SLA profile cannot be created, or the host cannot be placed in the existing SPI (the **host-limits** was exceeded), the static host definition will fail.

9.3.11 QoS for Subscribers and Hosts

9.3.11.1 QoS Parameters in Different Profiles

QoS aspects for subscribers and hosts can be defined statically on a SAP or dynamically using

Enhanced Subscriber Management. For example, in a VLAN-per-service model, different services belonging to a single subscriber are split over different SAPs, and thus the overall QoS (such as a scheduler policy) of this subscriber must be assigned using Enhanced Subscriber Management.

QoS parameters are shared among the subscriber profile and SLA profile as follows:

- The subscriber profile refers to HQoS ingress and egress scheduler policies which define the overall treatment for hosts of this subscriber when queues are used, or policers managed by HQoS are used at egress. If the subscriber is using policers, the subscriber profile also refers to CFHP ingress and egress policer-control-policies which define the overall treatment for hosts of this subscriber.
- The SLA profile refers to specific queue/policer settings for each host (BTV, VoIP, PC) using SAP ingress and SAP egress QoS policies. The SLA profile can also refer to an egress HQoS scheduler policy which defines the scheduling from the queues of the related host.

The primary use of the subscriber profile is to define the ingress and egress scheduler policies/ policer-control-policies used to govern the aggregate SLA for all hosts associated with a subscriber. To be effective, the queues/policers defined in the SLA profile’s QoS policies will reference a scheduler/arbiter from the scheduler policy/policer-control-policy respectively as their parent.

9.3.11.2 QoS Policy Overrides

Generic QoS queue/policer parameters could be specified for the SAP in a QoS policy and overridden for some customers by queue/policer parameters defined in the SLA profile. This allows for a single SAP ingress and SAP egress QoS policy to be used for many subscribers, while providing individual subscriber parameters for queue/policer operation.

9.3.12 ESM Subscriber Hierarchical Traffic Control

ESM subscribers can make use of both queues and/or policers for both the ingress and egress traffic. The queue and policers are configured within SAP ingress/egress policies applied to the SLA profile. The policers (at egress only) queue can parent to different levels/cir-levels with different weights/cir weights of a virtual scheduler configured within a scheduler policy, and to an egress port scheduler configured in a port scheduler policy, to achieve hierarchical traffic control. The policers can parent to different levels with different weights of an arbiter configured within a policer control policies to achieve hierarchical traffic control.

9.3.12.1 Subscriber HQoS

Hierarchical QoS (HQoS) corresponds to scheduling bandwidth distribution to policers, queues and schedulers and is applied using scheduler policies at ingress and egress of the subscriber profile for a subscriber, and at egress in the SLA profile for a host, together with a port scheduler at both the port and Vport level.

Each scheduler policy can contain up to three tiers of schedulers with lower level schedulers being able to parent to higher level schedulers in the same scheduler policy.

Policers and queues can parent to any scheduler in their related scheduler policy hierarchy (except Vport at egress) and also at the egress to a port scheduler.

Schedulers can parent to any higher level scheduler in their related scheduler policy hierarchy and, at the egress to a port scheduler configured within the port or Vport. When an egress port scheduler is used, an aggregate rate limit can be applied at the subscriber profile and Vport levels instead of using a scheduler. To extend the hierarchy further at egress, a tier 1 scheduler within a scheduler policy can parent to any scheduler in a scheduler policy at a higher level.

The scheduling levels are comprised of:

- Ingress and egress queues
- Egress policers
- Egress SLA-profile schedulers
- Ingress and egress subscriber profile schedulers
- Egress Vport schedulers
- Port schedulers

The ingress hierarchical parenting relationship options are shown in Figure 81.

Figure 81 **Ingress Scheduling Hierarchy Options**

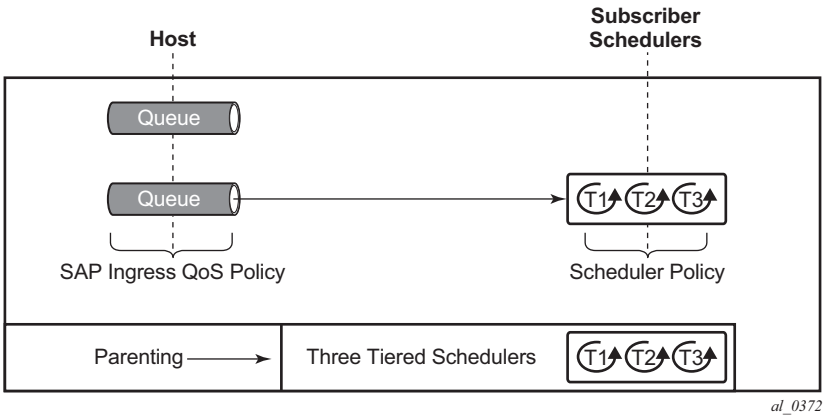

The egress hierarchical parenting relationship options are shown in Figure 82. Note that not all combinations can be configured concurrently, and some uses of port parent could be equally achieved using a scheduler parent and a child parent-location.

Figure 82 **Egress Scheduling Hierarchy Options**

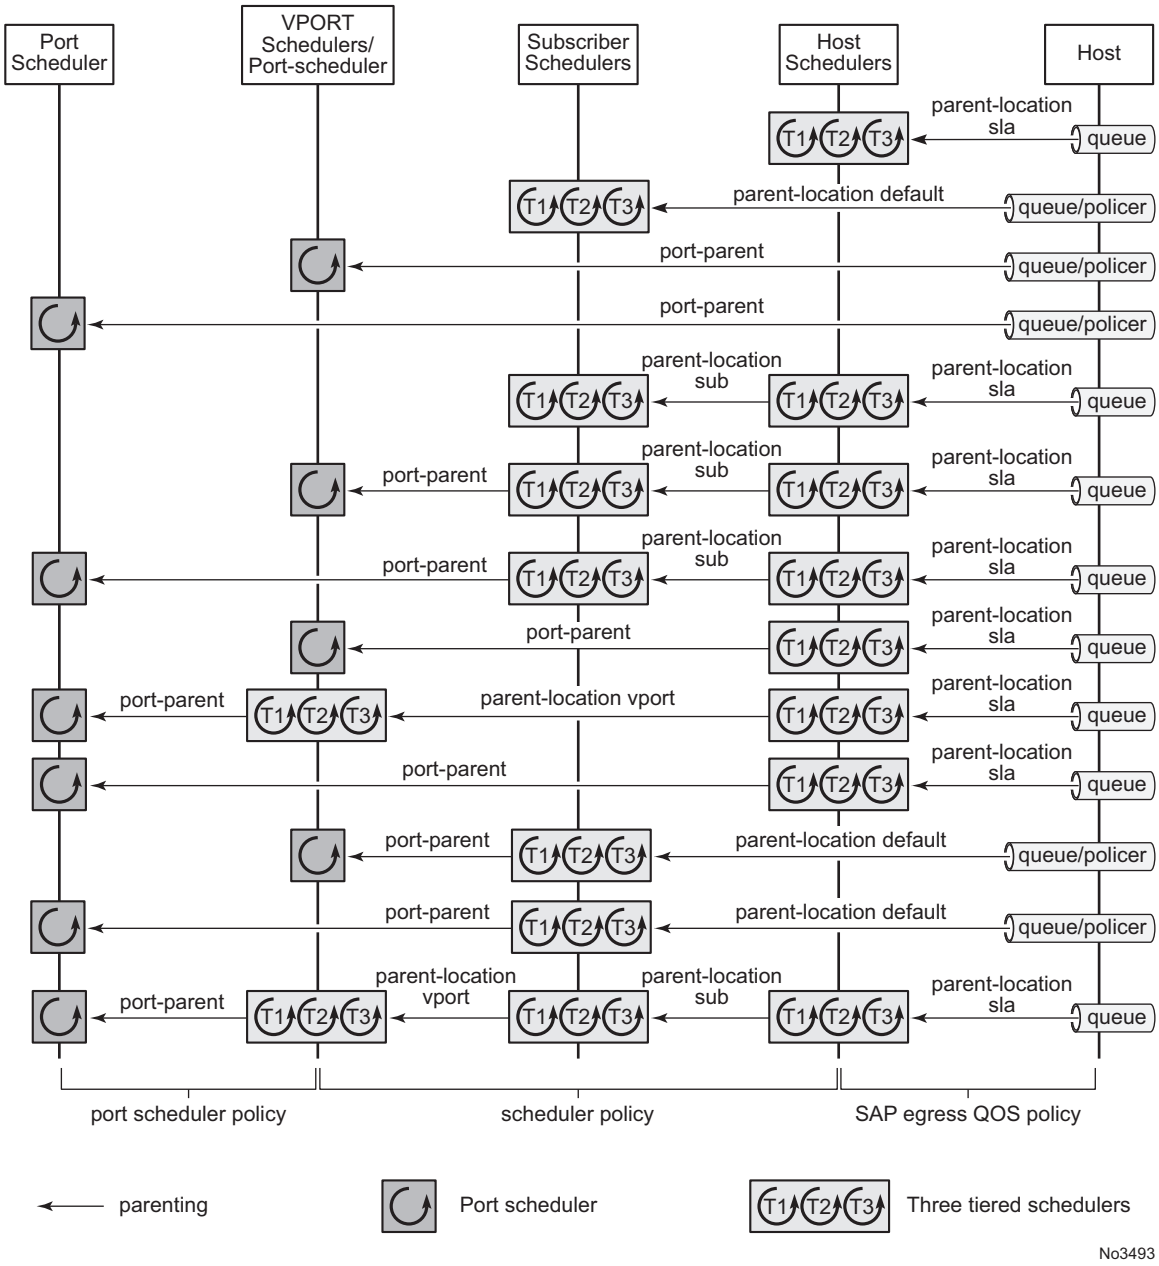

The **parent** command is used to specify the name of the parent scheduler when parenting a queue or scheduler, together with the level/cir-level and weight/cir-weight at which to connect.

```
config>qos>sap-ingress>queue# parent
  - parent scheduler-name [weight weight] [level level]
                          [cir-weight cir-weight] [cir-level cir-level]
```

```
config>qos>sap-egress>queue# parent
- parent scheduler-name [weight weight] [level level]
                        [cir-weight cir-weight] [cir-level cir-level]


config>qos>scheduler-policy>tier>scheduler# parent
- parent scheduler-name [weight weight] [level level]
                        [cir-weight cir-weight] [cir-level cir-level]
```

The location of the parent scheduler (in which applied scheduler policy it exists) for a policer or queue defaults to a scheduler in the subscriber ingress or egress scheduler policy. Parents of schedulers themselves must be explicitly configured and by default must be within the same scheduler policy.

At egress, the scheduler parenting relationship is determined using the **parent-location** command:

By default, egress queues parent to any scheduler in subscriber egress scheduler policy.

```
config>qos>sap-egress# parent-location default
```

Egress queues can parent to any scheduler within the scheduler policy applied to the egress of an SLA profile (this is not supported for policers managed by HQoS).

```
config>qos>sap-egress# parent-location sla
```

By default, a tier 1 scheduler in the scheduler policy is not allowed to be parented to another scheduler.

```
config>qos>scheduler-policy>tier# parent-location none
```

A tier 1 scheduler in the scheduler policy applied to the egress of an SLA profile can parent to a scheduler applied to the egress of a subscriber profile.

```
config>qos>scheduler-policy>tier# parent-location sub
```

A tier 1 scheduler in the scheduler policy applied to the egress of a subscriber profile can parent to a scheduler applied to the egress of a Vport.

```
config>qos>scheduler-policy>tier# parent-location vport
```

The configuration of a **parent-location** and frame-based accounting in a scheduler policy is mutually exclusive in order to ensure consistency between the different scheduling levels.

Note that the **parent-location** command is supported only on Ethernet interfaces. It is not supported for ESM over MPLS pseudowires.

Both egress queues and egress schedulers can port parent using directly to different levels/cir-levels, with different weights/cir weights, to a port egress port scheduler. Egress schedulers can also port parent directly to different levels/cir-levels, with different weights/cir weights, to a Vport egress port scheduler.

9.3.12.2 Subscriber CFHP

Class Fair Hierarchical Policing (CFHP) corresponds to the policing control of traffic by policers/arbiters. This uses policer control policies and can be applied for ingress and egress capacity control for the subscriber in the subscriber profile.

Each policer control policy can contain up to three tiers of arbiters with lower level arbiters being able to parent to higher level arbiters in the same scheduler policy.

Policers can parent to any arbiter in their related policer control policy hierarchy.

The policing levels are comprised of:

- Ingress and egress policers
- Ingress and egress subscriber arbiters

Note: Ingress policed traffic uses the shared policer-output-queues to access the switch fabric. At egress, the policed traffic accesses the egress port through a queue group queue (by default the policer-output-queues queue group, though user configurable queue groups can also be used) or a locally configured subscriber queue.

Egress policers can also be managed by HQoS.

The ingress hierarchical parenting relationship options are shown in Figure 83.

Figure 83 **Ingress Policing Hierarchy Options**

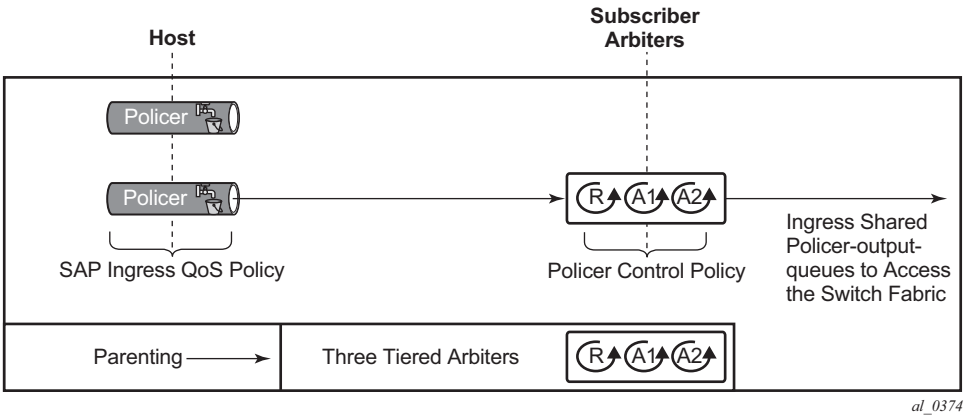

The egress hierarchical parenting relationship options are shown in Figure 84.

Figure 84 **Egress Policing Hierarchy Options**

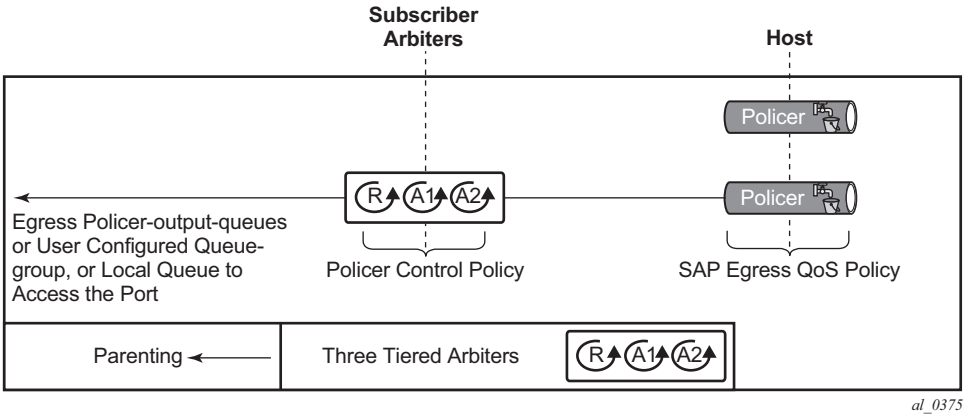

The **parent** command is used to specify the name of the parent arbiter when parenting a policer or arbiter, together with the level and weight at which to connect.

```
config>qos>sap-ingress>policer$ parent
  - parent arbiter-name [weight weight-level] [level level]

config>qos>sap-egress>policer$ parent
- parent arbiter-name [weight weight-level] [level level]

config>qos>plcr-ctrl-plcy>tier>arbiter# parent
- parent arbiter-name [weight weight-level] [level level]
```

9.3.12.3 ATM/Ethernet Last-Mile Aware QoS for Broadband Network Gateway

This feature allows the user to perform hierarchical scheduling of subscriber host packets such that the packet encapsulation overhead and ATM bandwidth expansion (when applicable) due to the last mile for each type of broadband session, that is, PPPoEoA LLC/SNAP and VC-Mux, IPoE, IPoEoA LLC/SNAP and VC-Mux, and so on, is accounted for by the 7450 ESS and 7750 SR acting as the Broadband Network Gateway (BNG).

The intent is that the BNG distributes bandwidth among the subscriber host sessions fairly by accounting for the encapsulation overhead and bandwidth expansion of the last mile such that packets are less likely to be dropped downstream in the DSLAM DSL port.

The last mile encapsulation type can be configured by the user or signaled using the Access-loop-encapsulation sub-TLV in the Vendor-Specific PPPoE Tags or DHCP Relay Options as per RFC 4679.

Furthermore, this feature allows the BNG to shape the aggregate rate of each subscriber and the aggregate rate of all subscribers destined to a given DSLAM to prevent congestion of the DSLAM. The subscriber aggregate rate is adjusted for the last mile overhead. The shaping to the aggregate rate of all subscribers of a given destination DSLAM is achieved via a new scheduling object, referred to as Virtual Port or Vport in CLI, which represents the DSLAM aggregation node in the BNG scheduling hierarchy

9.3.12.3.1 Broadband Network Gateway Application

An application of this feature in a BNG is shown in Figure 85.

Figure 85 **BNG Application**

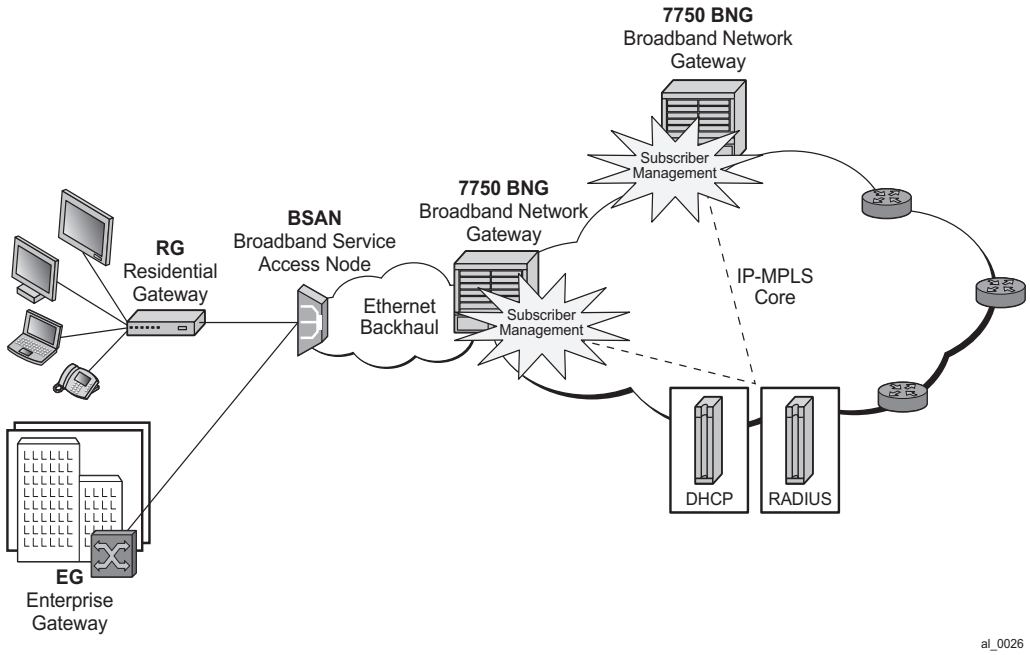

Residential and business subscribers use PPPoEoA, PPoA, IPoA, or IPoEoA based session over ATM/DSL lines. Each subscriber host can use a different type of session. Although Figure 1 illustrates ATM/DSL as the subscriber last mile, this feature supports both ATM and Ethernet in the last mile.

A subscriber SAP is auto-configured via DHCP or RADIUS authentication process, or is statically configured, and uses a Q-in-Q SAP with the inner C-VLAN identifying the subscriber while the outer S-VLAN identifies the Broadband Service Access Node (BSAN) which services the subscriber, such as, the DSLAM. The SAP configuration is triggered by the first successfully validated subscriber host requesting a session. Within each subscriber SAP, there can be one or more hosts using any of the above session types. The subscriber SAP terminates on an IES or VPRN service on the BNG. It can also terminate on a VPLS instance.

When the 7750 BNG forwards IP packets from the IP-MPLS core network downstream towards the Residential Gateway (RG) or the Enterprise Gateway (EG), it adds the required PPP and Ethernet headers, including the SAP encapsulation with C-VLAN/S-VLAN. When the BSAN node receives the packet, it strips the S-VLAN tag, strips or overwrites the C-VLAN tag, and adds padding to minimum Ethernet size if required. It also adds the LLC/SNAP or VC-mux headers plus the fixed AAL5 trailer and variable AAL5 padding (to next multiple of 48 bytes) and then segments the resulting PDU into ATM cells when the last mile is ATM/DSL. Thus the packet size will undergo a fixed offset due to the encapsulation change and a variable expansion

due to the AAL5 padding when applicable. Each type of subscriber host session will require a different amount of fixed offset and may require a per packet variable expansion depending of the encapsulation used by the session. The BNG node learns the encapsulation type of each subscriber host session by inspecting the Access-loop-encapsulation sub-TLV in the Vendor-Specific PPPoE Tags as specified in RFC 4679. The BNG node must account for this overhead when shaping packets destined to subscriber.

9.3.12.3.2 Queue Determination and Scheduling

Figure 86 illustrates the queuing and scheduling model for a BNG using the Ethernet or ATM last-mile aware QoS feature.

Figure 86 **BNG Queuing and Scheduling Model**

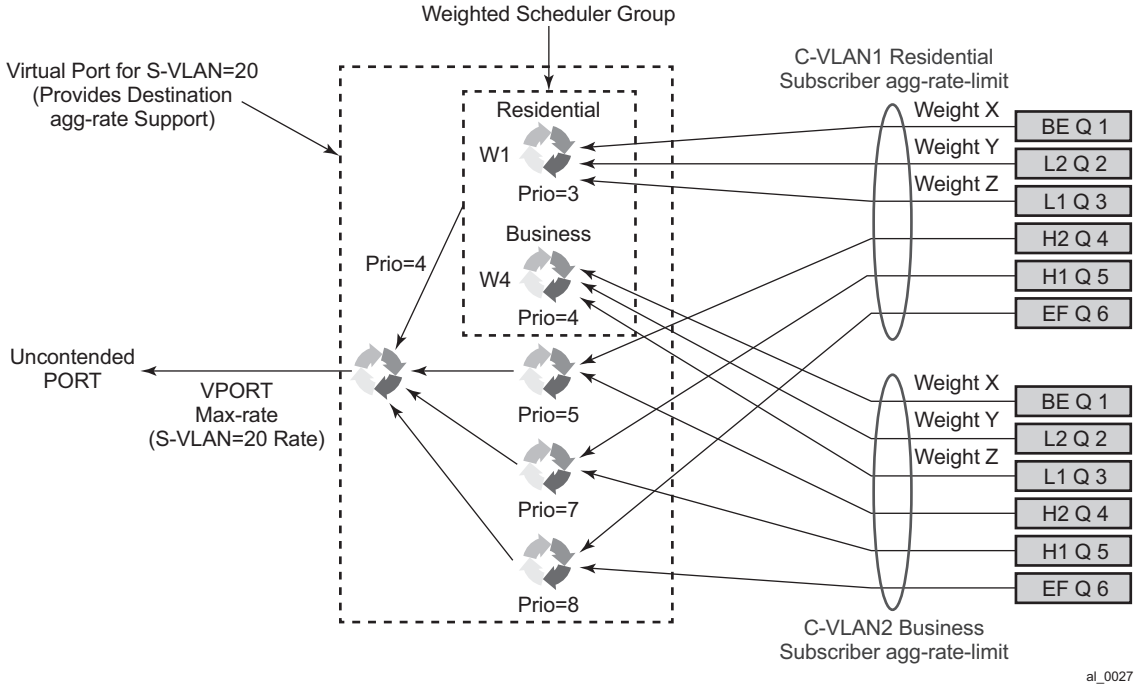

A set of per FC queues are applied to each subscriber host context to enforce the packet rate within each FC in the host session as specified in the subscriber's host SLA profile. A packet is stored in the queue corresponding the packet's FC as per the mapping of forwarding class to queue-id defined in the sap-egress QoS policy used by the host SLA profile. In the BNG application however, the host per FC queue packet rate is overridden by the rate provided in the RADIUS access-accept message. This rate represents the ATM rate that will be seen on the last mile, that is, it includes the encapsulation offset and the per packet expansion due to ATM segmentation into cells at the BSAN.

In order to enforce the aggregate rate of each destination BSAN, a scheduling node, referred to as virtual port, and Vport is in the CLI. The Vport operates exactly like a port scheduler with the difference that multiple Vport objects can be configured on the egress context of an Ethernet port. The user adds a Vport to an Ethernet port using the following command:

CLI Syntax: **config>port>ethernet>access>egress>vport** *'vport-name* **create**

The Vport is always configured at the port level even when a port is a member of a LAG. The *vport-name* is local to the port it is applied to but must be the same for all member ports of a LAG. It however does not need to be unique globally on a chassis.

CLI Syntax: **config>port>ethernet>access>egress>vport** *vport-name* **create**

The Vport is always configured at the port level even when a port is a member of a LAG. The *vport-name* is local to the port it is applied to but must be the same for all member ports of a LAG. It however does not need to be unique globally on a chassis.

The user applies a port scheduler policy to a Vport using the following command:

CLI Syntax: **config>port>ethernet>access>egress>vport>port-scheduler-policy** *port-scheduler-policy-name*

A Vport cannot be parented to the port scheduler when it is using a port scheduler policy itself. It is thus important the user ensures that the sum of the **max-rate** parameter value in the port scheduler policies of all Vport instances on a given egress Ethernet port does not oversubscribe the port's hardware rate. If it does, the scheduling behavior degenerates to that of the H/W scheduler on that port. A Vport which uses an agg-rate can be parented to a port scheduler. This is explained in Applying Aggregate Rate Limit to a Vport. Note that the application of the **agg-rate** *rate*, **port-scheduler-policy** and **scheduler-policy** commands under a Vport configuration are mutually exclusive.

Each subscriber host queue is port parented to the Vport which corresponds to the destination BSAN using the existing **port-parent** command:

CLI Syntax: `config>qos>sap-egress>queue>port-parent [weight weight] [level level] [cir-weight cir-weight] [cir-level cir-level]`

This command can parent the queue to either a port or to a Vport. These operations are mutually exclusive in CLI as explained above. When parenting to a Vport, the parent Vport for a subscriber host queue is not explicitly indicated in the above command. It is determined indirectly. The determination of the parent Vport for a given subscriber host queue is described in Vport Determination and Evaluation.

Furthermore, the weight (cir-weight) of a queue is normalized to the sum of the weights (cir-weights) of all active subscriber host queues port-parented at the same priority level of the Vport or the port scheduler policy. Since packets of ESM subscriber host queues are sprayed among the link of a LAG port based on the subscriber-id, it is required that all subscribers host queues mapping to the same Vport, such as having the same destination BSAN, be on the same LAG link so that the aggregate rate towards the BSAN is enforced. The only way of achieving this is to operate the LAG port in active/standby mode with a single active link and a single standby link.

The aggregate rate of each subscriber must also be enforced. The user achieves this by applying the existing **agg-rate-limit** command to the egress context of the subscriber profile:

CLI Syntax: `config>subscr-mgmt>sub-profile>egress>agg-rate-limit agg-rate`

In the BNG application however, this rate is overridden by the rate provided in the RADIUS access-accept message. This rate represents the ATM rate that will be seen on the last mile, that is, it includes the encapsulation offset and the per packet expansion due to ATM segmentation into cells at the BSAN.

9.3.12.3.3 Weighted Scheduler Group

The existing port scheduler policy defines a set of eight priority levels with no ability of grouping levels within a single priority. In order to allow for the application of a scheduling weight to groups of subscriber host queues competing at the same priority level of the port scheduler policy applied to the Vport, or to the Ethernet port, a new group object is defined under the port scheduler policy:

CLI Syntax: `config>qos>port-scheduler-policy>group group-name rate pir-rate [cir cir-rate]`

Up to eight groups can be defined within each port scheduler policy. One or more levels can map to the same group. A group has a rate and optionally a cir-rate and inherits the highest scheduling priority of its member levels. For example, the scheduler group shown in the Vport consists of level priority 3 and level priority 4. It thus inherits priority 4 when competing for bandwidth with the standalone priority levels 8, 7, and 5.

In essence, a group receives bandwidth from the port or from the Vport and distributes it within the member levels of the group according to the weight of each level within the group. Each priority level will compete for bandwidth within the group based on its weight under congestion situation. If there is no congestion, a priority level can achieve up to its rate (cir-rate) worth of bandwidth.

The mapping of a level to a group is performed as follows:

CLI Syntax: `config>qos>port-scheduler-policy>level priority-level rate pir-rate [cir cir-rate] group group-name [weight weight-in-group]`

Note: CLI will enforce that mapping of levels to a group are contiguous. In other words, a user would not be able to add priority level to group unless the resulting set of priority levels is contiguous.

When a level is not explicitly mapped to any group, it maps directly to the root of the port scheduler at its own priority like in existing behavior.

9.3.12.3.4 Queue and Subscriber Aggregate Rate Configuration and Adjustment

Software-Based Implementation (8.0R4)

The subscriber aggregate rate is adjusted and it will be based on an average frame size.

The user enables the use of this adjustment method by configuring the following option in the egress context of the subscriber profile:

CLI Syntax: `config>subscr-mgmt>sub-profile>egress>encap-offset [type type]`

This command allows the user to configure a default value to be used by all hosts of the subscriber in the absence of a valid signaled value. The following is a list of the configurable values:

Values: pppoa-llc, pppoa-null, pppoeoa-llc, pppoeoa-llc-fcs, pppoeoa-llc-tagged, pppoeoa-llc-tagged-fcs, pppoeoa-null, pppoeoa-null-fcs, pppoeoa-null-tagged, pppoeoa-null-tagged-fcs, ipoa-llc, ipoa-null, ipoeoa-llc, ipoeoa-llc-fcs, ipoeoa-llc-tagged, ipoeoa-llc-tagged-fcs, ipoeoa-null, ipoeoa-null-fcs, ipoeoa-null-tagged, ipoeoa-null-tagged-fcs, pppoe, pppoe-tagged, ipoe, ipoe-tagged

Otherwise, the fixed packet offset will be derived from the encapsulation type value signaled in the Access-loop-encapsulation sub-TLV in the Vendor-Specific PPPoE Tags as explained in Section Signaling of Last Mile Encapsulation Type. Only signaling using PPPoE Tags is supported in the software based implementation. The last signaled valid value is then applied to all active hosts of this subscriber. If no value is signaled in the subscriber host session or the value in the fields of the Access-loop-encapsulation sub-TLV are invalid, then the offset applied to the aggregate rate of this subscriber will use the last valid value signaled by a host of this subscriber if it exists, or the user entered default type value if configured, or no offset is applied.

The user also configures the average frame size value to be used for this adjustment:

CLI Syntax: `config>subscr-mgmt>sub-profile>egress>avg-frame-size` *bytes*

The value entered by the user must include the FCS but not the Inter-Frame Gap (IFG) or the preamble. If the user does not explicitly configure a value for the **avg-frame-size** parameter, then it will also be assumed the offset is zero regardless of the signaled or user-configured value.

The computation of the subscriber aggregate rate consists of taking the average frame size, adding the encapsulation fixed offset including the AAL5 trailer, and then adding the variable offset consisting of the AAL5 padding to next multiple of 48 bytes. The AverageFrameExpansionRatio is then derived as follows:

AverageFrameExpansionRatio = (53/48 x (AverageFrameSize + FixedEncapOffset + AAL5Padding)) / (AverageFrameSize + IFG + Preamble).

When the last mile is Ethernet, the formula simplifies to:

AverageFrameExpansionRatio = (AverageFrameSize + FixedEncapOffset + IFG + Preamble) / (AverageFrameSize + IFG + Preamble).

The following are the frame size and rate applied to the subscriber queue and scheduler:

Subscriber Host Queue (no change):

Size = ImmediateEgressEncap + Data

Rate = ImmediateEgressEncap + Data

Subscriber Aggregate Rate Scheduler:

Size = ImmediateEgressEncap + Data

Rate = sub-agg-rate / AverageFrameExpansionRatio

Note that the CPM applies the AverageFrameExpansionRatio adjustment to the various components used in the determination of the net subscriber operational aggregate rate. It then pushes these adjusted components to IOM which then makes the calculation of the net subscriber operational aggregate rate.

The formula used by the IOM for this determination is:

sub-oper-agg-rate = min(sub-policy-agg-rate/*AverageFrameExpansionRatio*, ancp_rate/*AverageFrameExpansionRatio*) + (igmp_rate_delta/ *AverageFrameExpansionRatio*),

where *sub-policy-agg-rate* is either the value configured in the **agg-rate-limit** parameter in the subscriber profile or the resulting RADIUS override value. In both cases, the CPM uses an internal override to download the adjusted value to IOM.

The value of *sub-oper-agg-rate* is stored in the IOM's subscriber table.

The following are the procedures for handling signaling changes or configuration changes affecting the subscriber profile:

1. If a new RADIUS update comes in for the aggregate subscriber rate, then a new subscriber aggregate ATM adjusted rate is computed by CPM using the last configured **avg-frame-size** and then programmed to IOM.
2. If the user changes the value of the **avg-frame-size** parameter, enables/ disables the **encap-offset** option, or changes the parameter value of the **encap-offset** option, the CPM will immediately trigger a re-evaluation of subscribers using the corresponding subscriber profile and an update the IOM with the new subscriber aggregate rate.
3. If the user changes the value of the **agg-rate-limit** parameter in a subscriber profile which has the **avg-frame-size** configured, this will immediately trigger a re-evaluation of subscribers using the corresponding subscriber profile. An update to the subscriber aggregate rate is performed for those subscribers which rate has not been previously overridden by RADIUS.

4. If the user changes the **type** value of the **encap-offset** command, this will immediately trigger a re-evaluation of subscribers using the corresponding subscriber profile. An update to the subscriber aggregate rate is performed for those subscribers which are currently using the default value.
5. If two hosts of the same subscriber signal two different encapsulation types, the last one signaled gets used at the next opportunity to re-evaluate the subscriber profile.
6. If a subscriber has a DHCP host, a static host or an ARP host, the subscriber aggregate rate will continue to use the user-configured default encapsulation type value or the last valid encapsulation value signaled in the PPPoE tags by other hosts of the same subscriber. If none was signaled or configured, then no rate adjustment is applied.

Hardware-Based Implementation

The data path will compute the adjusted frame size real-time for each serviced packet from a queue by adding the actual packet size to the fixed offset provided by CPM for this queue and variable AAL5 padding.

Like in the software based implementation, the user enables the use of the fixed offset and per packet variable expansion by configuring the following option in the egress context of the subscriber profile:

CLI Syntax: `config>subscr-mgmt>sub-profile>egress>encap-offset [type type]`

When this command is enabled, the fixed packet offset will be derived from the encapsulation type value signaled in the Access-loop-encapsulation sub-TLV in the Vendor-Specific PPPoE Tags or DHCP Relay Options as explained in Section Signaling of Last Mile Encapsulation Type.

If the user specifies an encapsulation type with the command, this value will be used as the default value for all hosts of this subscriber until a host session signaled a valid value. The signaled value is applied to this host only and the remaining hosts of this subscriber continue to use the user entered default type value if configured, or no offset is applied. Note however that hosts of the same subscriber using the same SLA profile and which are on the same SAP will share the same instance of FC queues. In this case, the last valid encapsulation value signaled by a host of that same instance of the SAP egress QoS policy will override any previous signaled or configured value.

The procedures for handling signaling changes or configuration changes affecting the subscriber profile are the same as in the Software based implementation with except for the following:

1. The **avg-frame-size** parameter in the subscriber profile is ignored.

2. If the user specifies an encapsulation type with the command, this value will be used as the default value for all hosts of this subscriber until a host session signaled a valid value. The signaled value is applied to this host and other hosts of the same subscriber sharing the same SLA profile and which are on the same SAP. The remaining hosts of this subscriber continue to use the user entered default type value if configured, or no offset is applied.
3. If the user enables/disables the **encap-offset** option, or changes the parameter value of the **encap-offset** option, the CPM will immediately trigger a re-evaluation of subscribers hosts using the corresponding subscriber profile and an update the IOM with the new fixed offset value.
4. If subscriber host session signals an encapsulation type at the session establishment time and subsequently sends a DHCP renewal message using a Layer 2 DHCP relay which does not insert option82 in a unicast message, the encapsulation type for this host will not change. Note that TR-101 states that option82 is mandatory for DHCP broadcast messages).
5. If a subscriber has a static host or an ARP host, the subscriber host will continue to use the user-configured default encapsulation type value or the last valid encapsulation value signaled in the PPPoE tags or DHCP relay options by other hosts of the same subscriber which use the same SLA profile instance. If none was signaled or configured, then no rate adjustment is applied.
6. The encapsulation type value signaled in DHCP relay options or PPPoE tags are not cross-checked against the host type. So, a host signaling PPPoA/LLC encapsulation type via DHCP relay options will be handled as if the packet included a PPPoE header when forwarded over the local Ethernet port. This results in applying an encap-offset in the data path which assumes the PPPoE header is added to forwarded packets over the local Ethernet port.

The **encap-offset** option forces all the rates to be either last-mile frame over the wire or local port frame over the wire, referred to as **LM-FoW** and **FoW** respectively. The system maintains a running average frame expansion ratio for each queue to convert queue rates between these two formats as explained in Frame Size, Rates, and Running Average Frame Expansion Ratio. Here are the details of the queue and scheduler operation:

1. When the **encap-offset** option is configured in the subscriber profile, the subscriber host queue rates, that is, CLI and operational PIR and CIR as well as queue bucket updates, the queue statistics, that is, forwarded, dropped, and HQoS offered counters use the **LM-FoW** format. The scheduler policy CLI and operational rates also use **LM-FoW** format. The port scheduler **max-rate** and the priority level rates and weights, if a Weighted Scheduler Group is used, are always entered in CLI and interpreted as **FoW** rates. The same is true for an **agg-rate-limit** applied to a Vport. Finally the subscriber **agg-rate-limit** is entered in CLI as **LM-FoW** rate. When converting between **LM-FoW** and **FoW** rates, the queue running average frame expansion ratio value is used.

- If the user enabled **frame-based-accounting** in a scheduler policy or **queue-frame-based-accounting** with subscriber **agg-rate-limit** and a port scheduler policy, the queue operational rate will be capped to a user configured **FoW** rate. The scheduler policy operational rates will also be in the **FoW** format. Note that a user configured queue **avg-frame-overhead** value is ignored since the running average frame expansion ratio is what is used when the **encap-offset** option is enabled.
- If the user configured queue **packet-byte-offset** value, it is ignored and is not accounted for in the net packet offset calculation.

2. When **no encap-offset** is configured in the subscriber profile, that is, default and pre-R9.0 behavior, queue CLI and operational PIR and CIR rates, as well as queue bucket updates, the queue statistics, use data format. The scheduler policy CLI and operational rates also use data format. The port scheduler **max-rate** and the priority level rates and weights, if a Weighted Scheduler Group is used, and the subscriber **agg-rate-limit** are entered in CLI and interpreted as **FoW** rates. When converting between **FoW** and data rates, the queue a**vg-frame-overhead** value is used and since this an Ethernet port, it is not user-configurable but constant and is equal to +20 bytes (IFG and preamble).
   - If the user enabled **frame-based-accounting** in a scheduler policy or **queue-frame-based-accounting** with subscriber **agg-rate-limit** and a port scheduler policy, the queue operational rate will be capped to a user configured FoW rate in CLI which is then converted into a data rate using the queue **avg-frame-overhead** constant value of +20 bytes. The scheduler policy operational rates will also be in the **FoW** format.
   - If the user configured queue **packet-byte-offset** value, it adjusts the immediate packet size. This means that the queue rates, that is, operational PIR and CIR, and queue bucket updates use the adjusted packet size. In addition, the queue statistics will also reflect the adjusted packet size. Scheduler policy rates, which are data rates, will use the adjusted packet size. The port scheduler **max-rate** and the priority level rates and weights, if a Weighted Scheduler Group is used, as well as the subscriber **agg-rate-limit** are always **FoW** rates and thus use the actual frame size. Packet byte offset settings are not included in the applied rate when (queue) frame based accounting is configured, however, the offsets are applied to the statistics.

9.3.12.3.5 Frame Size, Rates, and Running Average Frame Expansion Ratio

The following are the details of the rates and frame sizes applied to the subscriber host queues, the subscriber aggregate rate, and the Vport root scheduler for the scheduling model and when the **encap-offset** option is enabled in the subscriber profile.

Subscriber Host Queue:

Size = LastMileFrameOverWireEncap + Data

Rate = (48/53)* x (LastMileFrameOverWireEncap + Data)

*Applicable to ATM last-mile only.

Subscriber Aggregate Rate:

Size = LastMileFrameOverWireEncap + Data

Rate = (48/53)* x (LastMileFrameOverWireEncap + Data)

*Applicable to ATM last-mile only.

Vport/Port Port Scheduler and Weighted Scheduler Group

Size = FrameOverWireEncap + Data

Rate = FrameOverWireEncap + Data

When a frame arrives at the queue, it size will be *ImmediateEgressEncap+Data*. This size is stored as the *OfferedFrameSize* so that the queue offered stats used in HQoS calculations are correct. Refer to the HQoS-offered statistics as Offered.

This size is then adjusted by removing the *ImmediateEgressEncap* and adding the *LastMileFrameOverWireEncap*. This new adjusted frame size, referred as *LastMileOfferedFrameSize*, is then used for checking compliance of the frame against the queue PIR and CIR bucket sizes and for updating the queue forwarded and dropped stats.

The *LastMileOfferedFrameSize* value is computed dynamically for each packet serviced by the queue.

A new HQoS stat counter *OfferedLastMileAdjusted* is maintained for the purpose of calculating the running average frame expansion ratio, which is the ratio of the accumulated *OfferedLastMileAdjusted* and Offered of each queue:

RunningAverageFrameExpansionRatio = OfferedLastMileAdjusted / Offered

The **vport**/**port** *port-scheduler* will hand out its **FoW** bandwidth in terms of Fair Information Rate (FIR) bandwidth to each subscriber queue. This queue FIR must be converted into **LM-FoW** format to cap it by the queue PIR (*adminPIR*) and to make sure the sum of *FIRs* of all queues of the same subscriber does not exceed the subscriber **agg-rate-limit** which is also expressed in **LM-FoW** format. The conversion between these two rates makes use of the cumulative *RunningAverageFrameExpansionRatio* value.

Note that a queue **LM-foW** AdminPIR value will always be capped to the value of the local port **FoW** rate even if the conversion based on the current *RunningAverageFrameExpansionRatio* value indicates that a higher AdminPIR may be able to fill in the full line rate of the local port.

9.3.12.3.6 Vport Determination and Evaluation

In the BNG application, host queues of all subscribers destined to the same downstream BSAN, for example, all SAPs on the egress port matching the same S-VLAN tag value, are parented to the same Vport which matches the destination ID of the BSAN.

The BNG determines the parent Vport of a subscriber host queue, which has the **port-parent** option enabled, by matching the destination string **dest** string associated with the subscriber with the string defined under a Vport on the port associated with the subscriber.

The user configures the dest string match under the egress context of the Ethernet port associated with the subscriber:

CLI Syntax: **config>port>ethernet>access>egress>vport>host-match dest** *string* **create**

If a given subscriber host queue does not have the **port-parent** option enabled, it will be foster-parented to the Vport used by this subscriber and which is based on matching the **dest** string. If the subscriber could not be matched with a Vport on the egress port, the host queue will not be bandwidth controlled and will compete for bandwidth directly based on its own PIR and CIR parameters.

By default, a subscriber host queue with the **port-parent** option enabled is scheduled within the context of the port's port scheduler policy. In order to indicate the option to schedule the queue in the context of a port scheduler policy associated with a Vport, the user enters the following command in SLA profile used by the subscriber host:

CLI Syntax: **config>subscr-mgmt>sla-profile>egress>qos** *sap-egress-qos-policy-id* **vport-scheduler**

This command is persistent meaning that the user can re-enter the **qos** node without specifying the **vport-scheduler** argument each time and the system will remember it. The user can revert to the default setting without deleting the association of the SLA profile with the SAP egress QoS policy by explicitly re-entering the command with the following new argument:

CLI Syntax: **config>subscr-mgmt>sla-profile>egress>qos** *sap-egress-qos-policy-id* **port-scheduler**

9.3.12.3.7 Applying Aggregate Rate Limit to a Vport

The user can apply an aggregate rate limit to the Vport and apply a port scheduler policy to the port.

This model allows the user to oversubscribe the Ethernet port. The application of the **agg-rate** option is mutually exclusive with the application of a port scheduler policy, or a scheduler policy to a Vport.

When using this model, a subscriber host queue with the **port-parent** option enabled is scheduled within the context of the port's port scheduler policy. However, the user must still indicate to the system that the queues are managed by the aggregate rate limit instance of a Vport by enabling the **vport-scheduler** option in the subscriber host SLA profile:

CLI Syntax: **config>subscr-mgmt>sla-profile>egress>qos** *sap-egress-qos-policy-id* **vport-scheduler**

A subscriber host-queue which is port-parented will be parented to the port scheduler policy of the port used by the subscriber and aggregate rate limited within the instance of the Vport used by this subscriber and which is based on matching the **dest** string and **org** string. If the Vport exists but the port does not have a port scheduler policy applied, then the host queue will be orphaned and no aggregate rate limit can be enforced.

If a given subscriber host queue does not have the port-parent option enabled, it will be foster-parented to the port used by this subscriber and aggregate rate limited within the instance of the Vport used by this subscriber. If the Vport exists but the port does not have a port scheduler policy applied, then the host queue will be orphaned and no aggregate rate limit can be enforced.

9.3.12.3.8 Applying a Scheduler Policy to a Vport

The user can apply a scheduler policy to the Vport. This is allows scheduling control of subscriber tier 1 schedulers in a scheduler policy applied to the egress of a subscriber or SLA profile, or to a PW SAP in an IES or VPRN service.

The advantage of using a scheduler policy under a Vport, compared to the use of a port scheduler (with or without an **agg-rate** *rate*), is that it allows a port parent to be configured at the Vport level.

Bandwidth distribution from an egress port scheduler to a Vport configured with a scheduler policy can be performed based on the level/cir-level and weight/cir-weight configured under the scheduler's port parent. The result is in allowing multiple Vports, for example representing different DSLAMs, to share the port bandwidth capacity in a flexible way that is under the control of the user.

The configuration of a scheduler policy under a Vport is mutually exclusive with the configuration of a port scheduler policy or an aggregate rate limit.

A scheduler policy is configured under a Vport as follows:

CLI Syntax: **config>port>ethernet>access>egress>vport# scheduler-policy** *scheduler-policy-name*

When using this model, a tier 1 scheduler in a scheduling policy applied to a subscriber profile or SLA profiles must be configured as follows:

CLI Syntax: **config>qos>scheduler-policy>tier# parent-location vport**

If the Vport exists, but port does not have a scheduler policy applied, then its schedulers will be orphaned and no port level QOS control can be enforced.

The following **show**/**monitor**/**clear** commands are available related to the Vport scheduler:

```
show qos scheduler-hierarchy port port-id vport name [scheduler scheduler-name]
[detail]

show qos scheduler-stats port port-id vport name [scheduler scheduler-name] [detail]

monitor qos scheduler-stats port port-id vport name [interval seconds ] [repeat
repeat] [absolute|rate]

clear qos scheduler-stats port port-id vport name [scheduler scheduler-
name] [detail]
```

HQoS adjustment and host tracking are not supported on schedulers that are configured in a scheduler policy on a Vport, so the configuration of a scheduler policy under a Vport is mutually exclusive with the configuration of the egress-rate-modify parameter.

ESM over MPLS pseudowires are not supported when a scheduler policy is configured on a Vport.

9.3.12.3.9 Signaling of Last Mile Encapsulation Type

A subscriber host session can signal one of many encapsulation types each with a different fixed offset in the last mile. These encapsulation types are described in RFC 4679 and are illustrated in Figure 87 and Figure 88. The BNG node learns the encapsulation type of each subscriber host session by inspecting the Access-loop-encapsulation sub-TLV in the Vendor-Specific PPPoE Tags as specified in RFC 4679.When Ethernet is the last mile, the encapsulation type will result in a fixed offset for all packet sizes. When ATM/DSL is the last mile, there will be an additional expansion due to AAL5 padding to next multiple of 48 bytes and which varies depending on the packet size.

Both ATM and Ethernet access using PPP encapsulation options are supported in the software and hardware based implementations. Thus both provide support for the Access-loop-encapsulation sub-TLV in the Vendor-Specific PPPoEv4/PPPoEv6 Tags with the ATM encapsulation values and Ethernet encapsulation values.

ATM and Ethernet access using IP encapsulation are only supported using default encapsulation offset configuration in the subscriber profile in the software based implementation. Support for signaling the Access-loop-encapsulation sub-TLV in the DHCPv4/DHCPv6 Relay Options is included in the hardware based implementation. There is no support for DHCPv6 relay options.

Figure 87 **Subscriber Host Session Encapsulation Types**

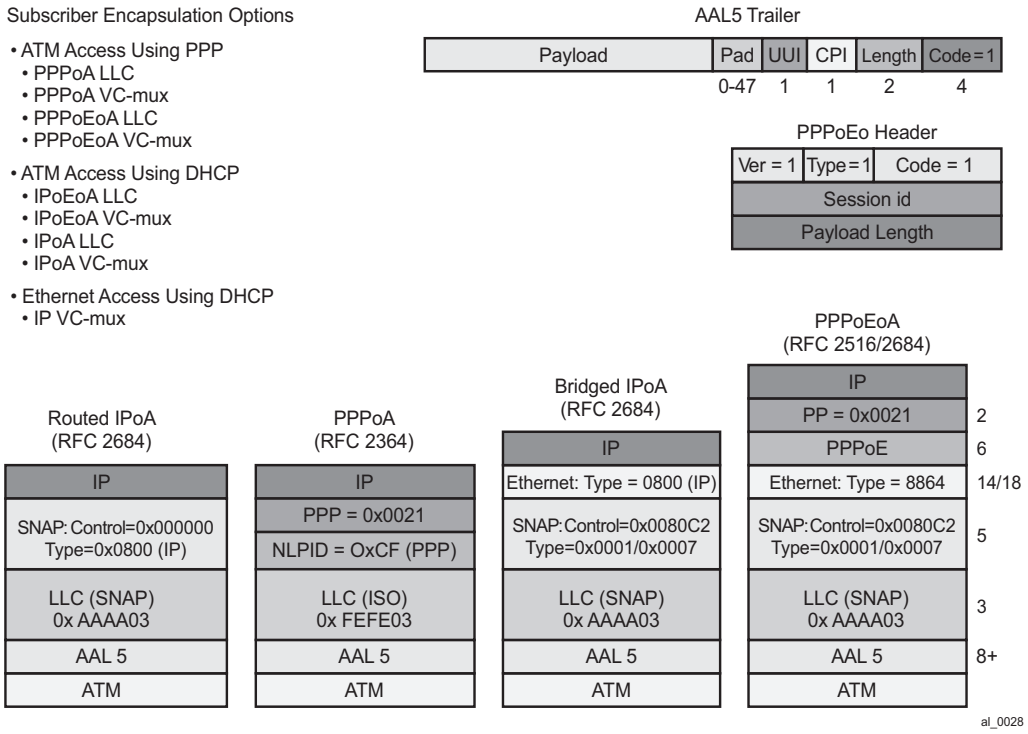

Figure 88 **Access-Loop-Encapsulation Sub-TLV**

Encapsulation combinations (RFC 4679)

```
                Access-loop-encapsulation
0                   1                   2
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|   Data Link   |    Encaps 1   |    Encaps 2   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

  Valid values for the sub-fields are as follows:

     Data Link

        0x00 AAL5
        0x01 Ethernet

     Encaps 1

        0x00 NA - Not Available
        0x01 Untagged Ethernet
        0x02 Single-Tagged Ethernet

     Encaps 2

        0x00 NA - Not Available
        0x01 PPPoA LLC
        0x02 PPPoA Null
        0x03 IPoA LLC
        0x04 IPoA Null
        0x05 Ethernet over AAL5 LLC with FCS
        0x06 Ethernet over AAL5 LLC without FCS
        0x07 Ethernet over AAL5 Null with FCS
        0x08 Ethernet over AAL5 Null without FCS
```

Encapsulation combinations

- AAL5
  - PPPoA LLC/Null
  - IPoA LLC/Null
  - Ethernet over ATM x 4
    - Tagged/Untagged PPP
    - Tagged/Untagged DHCP
  - Total of 20 AAL5combinations
- Ethernet
  - Tagged/Untagged PPP
  - Tagged/Untagged DHCP
  - Total of four Ethernet combinations
- *Total of 24 access combinations*

The operational last-mile values for hosts on the same sap, having the same SLA profile are displayed in following the show command:

CLI Syntax: **`show>service active-subscribers>ale-adjust`**

The data-link can have values: atm, other and unknown. If no offset is supplied it will be set to unknown. Other is used when the data-link is non-atm, otherwise it will state atm.

Operational per-queue values can also be found in the show command:

CLI Syntax: **`show>qos>scheduler-hierarchy`**

Here, one can see whether the queue is operating in last-mile mode. Note that this command is not available on HSMDAv2,

Last Mile ATM

```
*A:Dut-C# /show service active-subscribers ale-adjust
===============================================================================
Active Subscriber Access Loop Encapsulation adjustment
===============================================================================
Subscriber
```

```
   SAP                                              SLA profile
   Data-link Offset(bytes)
-------------------------------------------------------------------------------
hpolSub81
   1/1/11:2000.1                                    hpolSlaProf1
   atm       -10
-------------------------------------------------------------------------------
No. of Access Loop Encapsulation adjustments: 1
===============================================================================

*A:Dut-C# show qos scheduler-hierarchy subscriber "hpolSub81"
===============================================================================
Scheduler Hierarchy - Subscriber hpolSub81
===============================================================================
Ingress Scheduler Policy:
Egress Scheduler Policy :
-------------------------------------------------------------------------------
Root (Ing)
|
No Active Members Found on slot 1


Root (Egr)
| slot(1)
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->8->ATM  (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->7->ATM  (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->6->ATM  (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->5->ATM  (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->4->ATM  (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->3->ATM  (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->2->ATM  (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->1->ATM  (Port 1/1/11)
|
```

Last Mile Ethernet

```
*A:Dut-C# show service active-subscribers ale-adjust
===============================================================================
Active Subscriber Access Loop Encapsulation adjustment
===============================================================================
Subscriber
   SAP                                              SLA profile
   Data-link Offset(bytes)
-------------------------------------------------------------------------------
hpolSub81
   1/1/11:2000.1                                    hpolSlaProf1
   other     +12
-------------------------------------------------------------------------------
No. of Access Loop Encapsulation adjustments: 1
===============================================================================
```

```
*A:Dut-C# show qos scheduler-hierarchy subscriber "hpolSub81"
===============================================================================
Scheduler Hierarchy - Subscriber hpolSub81
===============================================================================
Ingress Scheduler Policy:
Egress Scheduler Policy :
-------------------------------------------------------------------------------
Root (Ing)
|
No Active Members Found on slot 1


Root (Egr)
| slot(1)
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->8->Eth  (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->7->Eth  (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->6->Eth  (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->5->Eth  (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->4->Eth  (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->3->Eth  (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->2->Eth  (Port 1/1/11)
|
|--(Q) : Sub=hpolSub81:hpolSlaProf1 2000->1/1/11:2000.1->1->Eth  (Port 1/1/11)
|
```

Next Mile ATM

```
*A:Dut-C# show service active-subscribers ale-adjust subscriber "hpolSub321"
===============================================================================
Active Subscriber Access Loop Encapsulation adjustment
===============================================================================
Subscriber
   SAP                                          SLA profile
   Data-link Offset(bytes)
-------------------------------------------------------------------------------
hpolSub321
   5/1/1:100/1                                  hpolSlaProf1
   atm       +8
-------------------------------------------------------------------------------
No. of Access Loop Encapsulation adjustments: 1
===============================================================================

*A:Dut-C# show qos scheduler-hierarchy subscriber "hpolSub321"
===============================================================================
Scheduler Hierarchy - Subscriber hpolSub321
===============================================================================
Ingress Scheduler Policy:
Egress Scheduler Policy :
-------------------------------------------------------------------------------
Root (Ing)
|
```

```
No Active Members Found on slot 5


Root (Egr)
| slot(5)
|--(Q) : Sub=hpolSub321:hpolSlaProf1 2000->5/1/1:100/1->8:ATM  (Port 5/1/1)
|
|--(Q) : Sub=hpolSub321:hpolSlaProf1 2000->5/1/1:100/1->7:ATM  (Port 5/1/1)
|
|--(Q) : Sub=hpolSub321:hpolSlaProf1 2000->5/1/1:100/1->6:ATM  (Port 5/1/1)
|
|--(Q) : Sub=hpolSub321:hpolSlaProf1 2000->5/1/1:100/1->5:ATM  (Port 5/1/1)
|
|--(Q) : Sub=hpolSub321:hpolSlaProf1 2000->5/1/1:100/1->4:ATM  (Port 5/1/1)
|
|--(Q) : Sub=hpolSub321:hpolSlaProf1 2000->5/1/1:100/1->3:ATM  (Port 5/1/1)
|
|--(Q) : Sub=hpolSub321:hpolSlaProf1 2000->5/1/1:100/1->2:ATM  (Port 5/1/1)
|
|--(Q) : Sub=hpolSub321:hpolSlaProf1 2000->5/1/1:100/1->1:ATM  (Port 5/1/1)
|
===============================================================================
```

9.3.12.3.10 Configuration Example

The following CLI configuration achieves the specific use case shown in Figure 86.

```
config
    qos
       port-scheduler-policy "dslam-vport-scheduler"
        group res-bus-be create
            rate 1000
        level 3 rate 1000 group res-bus-be weight w1
        level 4 rate 1000 group res-bus-be weight w4
        level 5 rate 1000 cir-rate 100
        level 7 rate 5000 cir-rate 5000
        level 8 rate 500 cir-rate 500
        max-rate 5000

       sap-egress 100                // residential policy
            queue 1                 // be-res
                port-parent weight x level 3
            queue 2                 // l2-res
                port-parent weight y level 3
            queue 3                 // l1-res
                port-parent weight z level 3
            queue 4                 // h2-res
                port-parent level 5
            queue 5                 // h1-res
                port-parent level 7
            queue 6                 // ef-res
                port-parent level 8
            fc be queue 1
            fc l2 queue 2
            fc l1 queue 3
            fc h2 queue 4
```

```
            fc h1 queue 5
            fc ef queue 6
        exit
        sap-egress 200                // business policy
            queue 1                 // be-bus
                     port-parent weight x level 4
            queue 2                 // l2-bus
                   port-parent weight y level 4
            queue 3                 // l1-bus
                      port-parent weight z level 4
            queue 4                 // h2-bus
                   port-parent level 5
            queue 5                 // h1-bus
                  port-parent level 7
            queue 6                 // ef-bus
                   port-parent level 8
            fc be queue 1
            fc l2 queue 2
            fc l1 queue 3
            fc h2 queue 4
            fc h1 queue 5
            fc ef queue 6
        exit
    exit

config
    sub-mgmt
        sla-profile "residential"
            egress
                qos 100 vport-scheduler
            exit
        exit
        sla-profile "business"
            egress
                qos 200 vport-scheduler
            exit
        exit
        sub-profile "residential"
            egress
               encap-offset
               avg-frame-size 1500
               agg-rate-limit 100
               exit
            exit
        exit
        sub-profile "business"
            egress
               encap-offset type pppoeoa-llc-tagged-fcs
               avg-frame-size 500
               agg-rate-limit 200
               exit
            exit
        exit
    exit

config
    port 1/1/1
        ethernet
            access
```

```
                egress
                    vport "dslam-1" create
                      port-scheduler-policy "dslam-vport-scheduler"
                        host-match dest "20" create
                        exit
                    exit
                exit
            exit
        exit
    exit
```

9.3.13 Subscriber Volume Statistics

Subscriber volume statistics or octet and packet counters are available through the queues and policers that are instantiated for the subscriber. The queue and policer configuration is defined in the SLA profile using ingress and egress QoS policy associations with optional overrides. By default, subscriber hosts that belong to the same subscriber, that are active on the same SAP, and that have the same SLA profile will share the set of queues and policers defined by that SPI. Alternatively, for bridged Residential Gateway scenarios, an SPI can be instantiated per subscriber session or per group identifier obtained during authentication. See SLA Profile Instance Sharing for more details. For HSMDA hardware, the ingress policer or queue and egress queue configuration is defined in the subscriber profile using ingress and egress QoS policy associations with optional overrides. All hosts belonging to the subscriber share the same set of queues and policers.

9.3.13.1 IP (Layer 3) Volume Accounting

Subscriber volume statistics by default count Layer 2 frame sizes optionally modified by configuration such as packet-byte-offset, last mile aware shaping, and so on.

To report subscriber volume statistics as Layer 3 (IP) packet sizes, the volume-stats-type can be configured to **ip** in the subscriber profile:

```
configure
    subscriber-mgmt
        sub-profile <subscriber-profile-name>
            volume-stats-type ip
```

volume-stats-type ip affects the subscriber statistics in SNMP, CLI, RADIUS accounting, XML accounting and Diameter Gx usage monitoring. Volume quota for RADIUS or Diameter Credit Control applications are interpreted as Layer 3 quota.

The following restrictions apply for **volume-stats-type ip**:

- Layer-3/IP accounting is not supported in combination with last-mile-aware shaping on HSMDA
- Layer-3/IP accounting is not supported in combination with ESMoPW on HSMDA
- Layer-3/IP accounting is not supported in combination with MLPPP
- Layer-3/IP accounting in combination with ESMoPW and last-mile-aware shaping may be inaccurate if the MPLS encapsulation overhead changes during the lifetime of a subscriber.
- Layer-3/IP accounting is restricted to a single encap per sla-profile instance (queue instance). The first host associated with the sla-profile instance (queue instance) determines the allowed encap. Conflicting encaps are:
  - PPPoE and IPoE on regular Ethernet SAPs
  - PPPoE and IPoE on PW SAPs
  - PPPoA and PPPoEoA on ATM ports
- PPPoE keep alive packets do not contain IP payload and introduce an error in Layer-3/IP accounting when enabled in combination with L2TP-LAC. A workaround is to isolate the keep alives in a separate queue/policer.
- Padding of frames smaller than the Ethernet minimum frame size (64B) may introduce an inaccuracy in Layer-3/IP accounting.
- With ATM in the last mile, last-mile-aware shaping may introduce an inaccuracy in Layer-3/IP accounting.
- Packet-Byte-Offset (PBO) changes during the lifetime of a subscriber introduces an inaccuracy in Layer-3/IP accounting.

9.3.13.2 Separate IPv4 and IPv6 Counters

IPv4 and IPv6 forwarded and dropped subscriber traffic can be counted separately via a **stat-mode v4-v6** command that is configured as a policer or queue qos override in the sla-profile. The **stat-mode v4-v6** command is only applicable for Enhanced Subscriber Management (ESM).

```
configure subscriber-mgmt
        sla-profile "sla-profile-1" create
            ingress
                qos 10
                    queue 1
                        stat-mode v4-v6
                    exit
                    policer 1
                        stat-mode v4-v6
                    exit
```

```
            exit
        exit
        egress
            qos 10
                queue 1
                    stat-mode v4-v6
                exit
                policer 1
                    stat-mode v4-v6
                exit
            exit
        exit
    exit
```

For policers, the **stat-mode** command overrides the policer stat-mode configuration as defined in the sap-ingress or sap-egress qos policy. For details on sap-ingress and sap-egress policer stat-mode, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Quality of Service Guide*. For a policer in stat-mode v4-v6, following counters are available:

- Offered IPv4 octets and packets
- Offered IPv6 octets and packets
- Dropped IPv4 octets and packets
- Dropped IPv6 octets and packets
- Forwarded IPv4 octets and packets
- Forwarded IPv6 octets and packets

When a policer's stat-mode is changed while the sla profile is in use, any previous counter values are lost and any new counters are set to zero.

For queues, a stat-mode is only available for use in Enhanced Subscriber Management (ESM) context to enable separate IPv4/IPv6 counters. For a queue in stat-mode v4-v6, following counters are available:

- Offered High Priority, Low Priority, Uncolored, Managed octets and packets
- Dropped IPv4 octets and packets
- Dropped IPv6 octets and packets
- Forwarded IPv4 octets and packets
- Forwarded IPv6 octets and packets

A queue's stat-mode cannot be changed while the sla profile is in use.

Note that there will be no in-profile/out-of-profile forwarded and dropped counters for policers and queues in **stat-mode v4-v6**.

Non IP traffic (for example PPPoE LCP frames) is counted against the IPv4 counters.

The separate IPv4 and IPv6 forwarded and dropped counters are reported in

- SNMP
- CLI

```
show service active-subscribers detail
- - - snip - - -
-------------------------------------------------------------------------------
SLA Profile Instance statistics
-------------------------------------------------------------------------------
                        Packets                 Octets

Off. HiPrio           : 0                       0
Off. LowPrio          : 1102685                 1102685000
Off. Uncolor          : 0                       0
Off. Managed          : 0                       0

Queueing Stats (Ingress QoS Policy 10)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
Dro. V4               : 0                       0
Dro. V6               : 0                       0
For. V4               : 367543                  367543000
For. V6               : 735142                  735142000

Queueing Stats (Egress QoS Policy 10)
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
Dro. V4               : 0                       0
Dro. V6               : 0                       0
For. V4               : 367543                  367543000
For. V6               : 735088                  735088000

-------------------------------------------------------------------------------
SLA Profile Instance per Queue statistics
-------------------------------------------------------------------------------
                        Packets                 Octets

Ingress Queue 1 (Unicast) (Priority) (Stats mode: v4-v6)
Off. HiPrio           : 0                       0
Off. LowPrio          : 1102685                 1102685000
Dro. V4               : 0                       0
Dro. V6               : 0                       0
For. V4               : 367545                  367545000
For. V6               : 735146                  735146000

Egress Queue 1 (Stats mode: v4-v6)
Dro. V4               : 0                       0
Dro. V6               : 0                       0
For. V4               : 367547                  367547000
For. V6               : 735096                  735096000

-------------------------------------------------------------------------------
SLA Profile Instance per Policer statistics
```

```
-------------------------------------------------------------------------------
                          Packets                  Octets

Ingress Policer 1 (Stats mode: v4-v6)
Off. V4                : 0                        0
Off. V6                : 0                        0
Dro. V4                : 0                        0
Dro. V6                : 0                        0
For. V4                : 0                        0
For. V6                : 0                        0

Egress Policer 1 (Stats mode: v4-v6)
Off. V4                : 0                        0
Off. V6                : 0                        0
Dro. V4                : 0                        0
Dro. V6                : 0                        0
For. V4                : 0                        0
For. V6                : 0                        0
```

- RADIUS accounting

When a queue or policer is configured in stat-mode v4-v6, existing VSA's are re-used in RADIUS detailed per queue / per policer accounting (**configure subscriber-mgmt radius-accounting-policy** <*name*> **include-radius-attribute detailed-acct-attributes**):

- in-profile counter VSA's map to IPv4 octets/packets
- ingress queue high priority dropped counter VSA's map to IPv4 octets/packets
- out-of-profile counter VSA's map to IPv6 octets/packets
- ingress queue low priority dropped counter VSA's map to IPv6 octets/packets

In addition the [26-6527-107] Alc-Acct-I-statmode / [26-6527-127] Alc-Acct-O-statmode is sent with value set to "v4-v6".

Optionally a set of VSA's can be included in RADIUS accounting to report the aggregate IPv6 forwarded octets and packets of queues and policers with stat-mode v4-v6 enabled (configure subscriber-mgmt radius-accounting-policy <name> include-radius-attribute detailed-acct-attributes v6-aggregate-stats):

[26-6527-194] Alc-IPv6-Acct-Input-Packets

[26-6527-195] Alc-IPv6-Acct-Input-Octets

[26-6527-196] Alc-IPv6-Acct-Input-GigaWords

[26-6527-197] Alc-IPv6-Acct-Output-Packets

[26-6527-198] Alc-IPv6-Acct-Output-Octets

[26-6527-199] Alc-IPv6-Acct-Output-Gigawords

Refer to the *7750 SR and VSR RADIUS Attributes Reference Guide* for a detailed description of all counter attributes.

- XML accounting

The complete-subscriber-ingress-egress and custom-record-subscriber XML records use following fields to represent IPv4 and IPv6 forwarded/dropped octets and packets for queues or policers with sta**t-mode v4-v**6 enabled:

v4po - IPv4PktsOffered (policer only)

v4oo - IPv4OctetsOffered (policer only)

v6po - IPv6PktsOffered (policer only)

v6oo - IPv6OctetsOffered (policer only)

v4pf - IPv4PktsForwarded

v6pf - IPv6PktsForwarded

v4pd - IPv4PktsDropped

v6pd - IPv4PktsDropped

v4of - IPv4OctetsForwarded

v6of - IPv6OctetsForwarded

v4od - IPv4OctetsDropped

v6od - IPv4OctetsDropped

For custom records, the following CLI is re-used to include v4/v6 counters if the queue is configured in **stat-mode v4-v6**:

```
i-counters
        all-packets-offered-count          # n/a
        all-octets-offered-count           # n/a
        high-packets-offered-count         # n/a
        low-packets-offered-count          # n/a
        uncoloured-packets-offered-count   # n/a
```

```
            high-octets-offered-count          # n/a
            low-octets-offered-count           # n/a
            uncoloured-octets-offered-count    # n/a
            all-packets-offered-count          # n/a
            all-octets-offered-count           # n/a
            high-packets-discarded-count       # IPv4
            low-packets-discarded-count        # IPv6
            high-octets-discarded-count        # IPv4
            low-octets-discarded-count         # IPv6
            in-profile-packets-forwarded-count  # IPv4
            out-profile-packets-forwarded-count # IPv6
            in-profile-octets-forwarded-count   # IPv4
            out-profile-octets-forwarded-count  # IPv6

        e-counters
            in-profile-packets-forwarded-count  # IPv4
            in-profile-packets-discarded-count  # IPv4
            out-profile-packets-forwarded-count # IPv6
            out-profile-packets-discarded-count # IPv6
            in-profile-octets-forwarded-count   # IPv4
            in-profile-octets-discarded-count   # IPv4
            out-profile-octets-forwarded-count  # IPv6
            out-profile-octets-discarded-count  # IPv6
```

On HSMDA, **stat-mode v4-v6** is configured as a policer or queue qos override in the subscriber profile:

```
configure subscriber-mgmt
        sub-profile "sub-profile-2" create
            hsmda
                egress-qos
                    qos 10
                        queue 1
                            stat-mode v4-v6
                        exit
                    exit
                exit
                ingress-qos
                    qos 10
                        policer 1
                            stat-mode v4-v6
                        exit
                        queue 1
                            stat-mode v4-v6
                        exit
                    exit
                exit
            exit
        exit
```

The stat-mode on egress hsmda queues is always enabled per subscriber: when enabling **stat-mode v4-v6** for one hsmda queue, it is automatically enabled for all hsmda queues for that subscriber profile.

9.3.14 Configuring IP and IPv6 Filter Policies for Subscriber Hosts

Access Control Lists (ACLs) for subscriber traffic are defined as IP and IPv6 filter policies and are configured in the SLA-profile associated with the subscriber. For information about IP and IPv6 filter policy configurations, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Router Configuration Guide*.

```
config>subscr-mgmt>sla-prof
        sla-profile sla-profile-1 create
                ingress
                        ip-filter 100
                        ipv6-filter 300
                exit
                egress
                        ip-filter 200
                        ipv6-filter 400
                exit
            exit
```

Traffic from different hosts of a single subscriber and associated with the same sla-profile instance, is subject to the filter policies defined in the SLA profile.

The IP or IPv6 filter policy configuration of subscriber hosts can be dynamically updated using different mechanisms:

1. Assign a new SLA profile.

   This can be done dynamically by, for example, a RADIUS CoA message. As the SLA profile also defines the QoS configuration for the subscriber hosts, this change may result in a discontinuity in accounting.

   Changing the ip-filter policy in an SLA profile in use by an active subscriber is allowed in the CLI, but not recommended. Changing the IPv6 filter policy in an SLA profile in use by an active subscriber is prevented in the CLI.

2. Override the IP and IPv6 filter policies

   Alternatively, it is also possible to dynamically override the IP and IPv6 filter policies per subscriber-host through a RADIUS Access-Accept or CoA message. Following VSA should be included in the RADIUS message:

3. Insert subscriber host specific filter entries

   A subscriber host specific entry is dynamically created from a RADIUS access accept or CoA message using the NAS-Filter-Rule or Alc-Ascend-Data-Filter-Host-Spec attribute (see also IP Filter Attribute Format Details for a detailed description of the attribute format):

Table 77 **NAS-Filter-Rule or Alc-Ascend-Data-Filter-Host-Spec Attributes**

Attribute ID	Attribute Name	Type	Limits	SR OS Format
92	NAS-Filter-Rule	String	max. 10 attributes per message or max. 10 filter entries per message	The format of a NAS-Filter-Rule is defined in RFC-3588 section-4.3. A single filter rule is a string of format “<action> <direction> <protocol> from <source> to <destination> <options>”. Multiple rules should be separated by a NUL (0x00). A NAS-Filter-Rule attribute may contain a partial rule, one rule, or more than one rule. Filter rules may be continued across attribute boundaries. A RADIUS message with NAS-Filter-Rule attribute value equal to 0x00 or “ “ (a space) removes all host specific filter entries for that host. See also IP Filter Attribute Format Details. For example: Nas-Filter-Rule = "permit in ip from any to 10.1.1.1/32"

Table 77 **NAS-Filter-Rule or Alc-Ascend-Data-Filter-Host-Spec Attributes (Continued)**

Attribute ID	Attribute Name	Type	Limits	SR OS Format
26-6527-159	Alc-Ascend-Data-Filter-Host-Spec	octets	max. 10 attributes per message or max. 10 filter entries per message. min. length 22 bytes (IPv4), 46 bytes (IPv6) max. length: 110 bytes (IPv4), 140 bytes (IPv6)	A string of octets with fixed field length (type (ipv4/ipv6), direction (ingress/egress), src-ip, dst-ip, and so on). Each attribute represents a single filter entry. See IP Filter Attribute Format Details for a description of the format. For example: # "permit in ip from any to 10.1.1.1/32" Alc-Ascend-Data-Filter-Host-Spec = 0x0101010000000000a010101002000000000000000000

The Alc-Subscriber-Filter VSA is a comma separated list of strings:

Table 78 **Alc-Subscriber-Filter VSA Fields**

Field	Use
Ingr-v4:<number>	Ingress ipv4 filter
Egr-v4:<number>	Egress ipv4 filter
Ingr-v6:<number>	Ingress ipv6 filter
Egr-v6:<number>	Egress ipv6 filter

The filter number can have following values:

Table 79 **Filter Number Results**

<number>	Result
1..65535	Ignore filter from sla-profile configuration and assign corresponding pre-configured filter
0	Ignore filter from sla-profile configuration and do not assign a new filter (only allowed if no dynamic subscriber host specific rules are present)
-1	No change in filter configuration
-2	Restore filter from sla-profile configuration

Notes:

- Irrelevant fields (IPv4 filters for an IPv6 host) will be ignored.
- RADIUS CoA message: if the ingress or egress field is missing in the VSA, there will be no change for that direction.
- RADIUS Access-Accept message: if the ingress or egress field is missing in the VSA, then the IP filters as specified in the SLA profile will be active for that direction.

An SLA profile IP filter override is applicable to all dynamic host types, including L2TP LNS but excluding L2TP LAC.

- Insert subscriber host specific filter entries

  A subscriber host specific entry is a filter entry where the match criteria is automatically extended with the subscriber host IP or IPv6 address as source (ingress) or destination (egress) IP. They represent a per host customization of a generic filter policy: only traffic to/from the subscriber host will match against these entries.

  A subscriber host specific entry is dynamically created from a RADIUS access accept or CoA message using the NAS-Filter-Rule attribute:

The format used to specify host specific filter entries (NAS-Filter-Rule format or Alc-Ascend-Data-Filter-Host-Spec format) cannot change during the lifetime of the subscriber host. A RADIUS message can only contain a single format for host specific filter entries.

Up to 10 host specific filter rules can be specified in a single RADIUS message. Each new RADIUS CoA message containing host specific filter attributes overwrites the previous subscriber host-specific filter entries for that host provided that there are enough free entries in the reserved range.

Subscriber host specific filter entries can be removed with a RADIUS CoA message with NAS-Filter-Rule attribute value equal to 0x00 or “ “ (a space).

When the subscriber host session terminates or is disconnected the corresponding subscriber host specific filter entries are also deleted.

Note that subscriber host-specific filter entries are moved if the subscriber host filter policy is changed (new SLA profile or ip filter policy override) and the new filter policy contains enough free reserved entries (sub-insert-RADIUS).

A range of entries must be reserved for subscriber host specific entries in a filter policy:

```
config>filter
        ip-filter 100 create
```

```
                                sub-insert-radius start-entry 1000 count
                        100
```

High and low watermarks can be configured to raise an event when the thresholds of free entries in the reserved range are reached:

```
config>filter>ip-filter# sub-insert-wmark ?
  - no sub-insert-wmark
  - sub-insert-wmark low <low-watermark> high <high-
 watermark>
```

<low-watermark> : [0..100]

<high-watermark> : [0..100]

Use following show commands to check filter policy details and the filter configuration for a subscriber host:

```
# show filter ip <ip-filter-id> type <entry-type>
# show filter ipv6 <ipv6-filter-id> type <entry-type>
     <entry-type>        : fixed|radius-insert|credit-control-insert

# show service active-subscribers filter [subscriber <sub-ident-string>] [origin
<origin>]
     <sub-ident-string>   : [32 chars max]
     <origin>             : radius|credit-control
```

- Insert shared filter entries

  The target application for RADIUS shared filter entries is operators that have a predefined limited number of different filter lists that each are shared with multiple subscriber hosts and that are to be managed and activated from RADIUS at authentication.

  A local configured ip or ipv6 filter associated with a host (sla-profile or host filter override) can be enhanced with dynamic filter entries that can be shared with multiple subscriber hosts. The shared dynamic filter entries are inserted with a set of RADIUS attributes "[242] Ascend-Data-Filter" or "[26-6527-158] Alc-Nas-Filter-Rule-Shared" received in a RADIUS Access-Accept or CoA message. A CoA message containing a set of one of those attributes overrides the previous set of shared filter entries active for that subscriber host.

For each unique set of dynamic filter entries received per type (ipv4/ipv6) and direction (ingress/egress), a copy is made of the local filter with the dynamic entries included at a preconfigured insert point. If the same set of dynamic filter entries is sent to subscriber hosts that have the same associated local filter, then they will share the same filter copy. When there are no more subscriber hosts associated with a filter copy, then the filter copy is deleted. A filter copy is identified as *local filter id*:*number*. For example: show filter ip 10:2

Shared filter entries are moved if the subscriber host filter policy is changed (new SLA profile or ip filter policy override) and if the new filter policy contains enough free reserved entries.

Figure 89 **Insert Shared Filters**

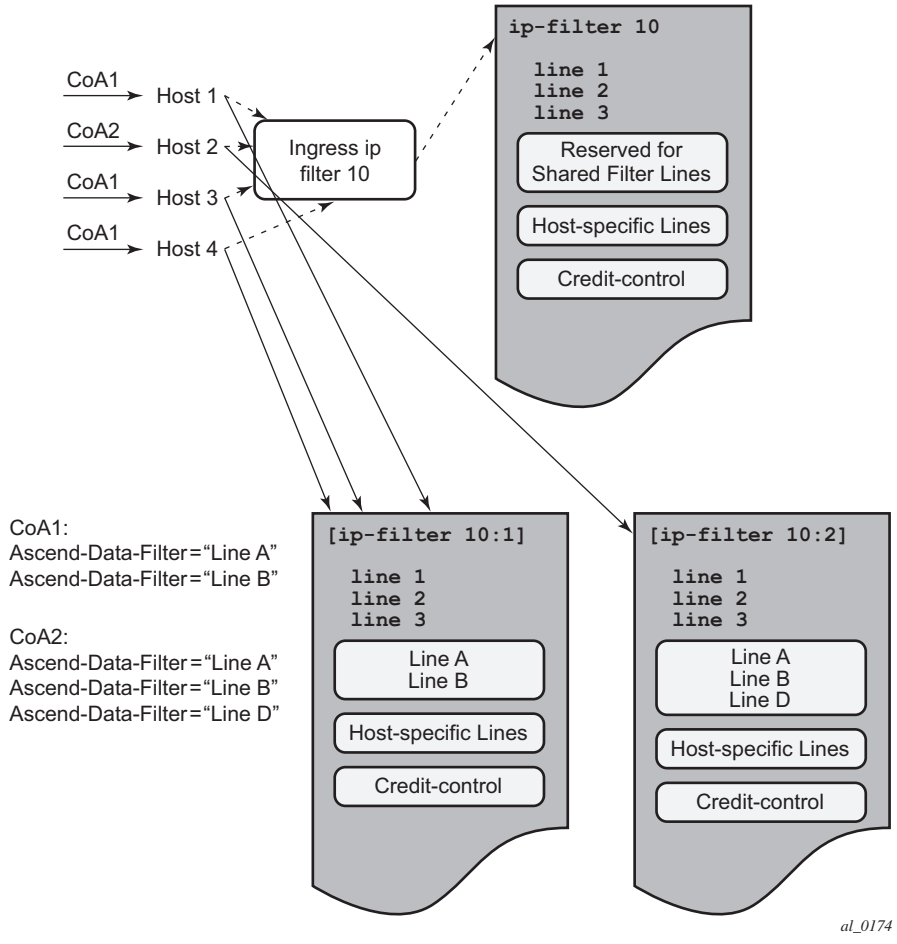

A range of entries must be reserved for shared entries in a filter policy:

```
config>filter
ip-filter 10 create
```

```
 sub-insert-shared-radius start-entry 100 count
10
```

High and low watermarks can be configured to raise an event when the thresholds of dynamic filter copies are reached:

```
config>filter>ip-filter# shared-radius-filter-wmark ?
        - no shared-radius-filter-wmark
        - shared-radius-filter-wmark low low-watermark
        high high-watermark
               low-watermark  : [0..8000]
               high-watermark : [0..8000]
```

The format used to specify shared filter entries (Alc-Nas-Filter-Rule-Shared format or Ascend-Data-Filter format) cannot change during the lifetime of the subscriber host. A RADIUS message can only contain a single format for shared filter entries.

Shared filter entries can be removed with a RADIUS CoA message with Alc-Nas-Filter-Rule-Shared attribute value equal to 0x00 or “ “ (a space).

Table 80 **Alc-Nas-Filter-Rule-Shared Attributes**

Attribute ID	Attribute Name	Type	Limits	SR OS Format
242	Ascend-Data-Filter	Octets	multiple attributes per RADIUS message allowed. min. length 22 bytes (IPv4), 46 bytes (IPv6) max. length: 110 bytes (IPv4), 140 bytes (IPv6)	a string of octets with fixed field length (type (ipv4/ipv6), direction (ingress/egress), src-ip, dst-ip, ...). Each attribute represents a single filter entry. See IP Filter Attribute Format Details for a description of the format. For example: # "permit in ip from any to 10.1.1.1/32" Ascend-Data-Filter = 0x01010100000000000a010101002000000000000000000

Table 80 **Alc-Nas-Filter-Rule-Shared Attributes (Continued)**

Attribute ID	Attribute Name	Type	Limits	SR OS Format
26-6527-158	Alc-Nas-Filter-Rule-Shared	string	Multiple attributes per RADIUS message allowed.	The format is identical to [92] NAS-Filter-Rule and is defined in rfc-3588 section-4.3. A single filter rule is a string of format “<action> <direction> <protocol> from <source> to <destination> <options>”. Multiple rules should be separated by a NUL (0x00). An Alc-Nas-Filter-Rule-Shared attribute may contain a partial rule, one rule, or more than one rule. Filter rules may be continued across attribute boundaries. A RADIUS message with Alc-Nas-Filter-Rule-Shared attribute value equal to 0x00 or “ “ (a space) removes the shared filter entries for that host. See also IP Filter Attribute Format Details. For example: Alc-Nas-Filter-Rule-Shared = "permit in ip from any to 10.1.1.1/32"

9.3.14.1 IP Filter Attribute Format Details

The format for [92] Nas-Filter-Rule and [26-6527-158] Alc-Nas-Filter-Rule-Shared is a string formatted as *action direction protocol* from *source* to *destination options*. Table 81 provides details on the respective fields.

Table 81 **Nas-Filter-Rule and Alc-Nas-Filter-Rule-Shared Fields**

Action or Classifier	Value	Corresponding SR OS Filter Function
<direction>	in	ingress
	out	egress

Table 81 **Nas-Filter-Rule and Alc-Nas-Filter-Rule-Shared Fields (Continued)**

Action or Classifier	Value		Corresponding SR OS Filter Function
<protocol>	ip		—
	any number [0 to 255]		—
	ip		—
	any number [1 to 42]		—
	any number [45 to 49]		—
	any number [52 to 59]		—
	any number [61 to 255]		—
	any number 43|44|50|51|60		—
from <source>	any	100	ingress: src-ip = host-ip-address; src-port eq 100 egress: src-ip = 0.0.0.0/0 | ::/0; src-port eq 100
		200 to 65535	ingress: src-ip = host-ip-address; src-port range 200 65535 egress: src-ip = 0.0.0.0/0 | ::/0; src-port range 200 65535
	ip-prefix/ length	100	ingress: src-ip = host-ip-address; src-port eq 100 egress: src-ip = ip-prefix/length; src-port eq 100
		200 to 65535	ingress: src-ip = host-ip-address; src-port range 200 65535 egress: src-ip = ip-prefix/length; src-port range 200 65535
	any	100	ingress: dst-ip = 0.0.0.0/0 | ::/0; dst-port eq 100 egress: dst-ip = host-ip-address; dst-port eq 100
		200 to 65535	ingress: dst-ip = 0.0.0.0/0 | ::/0; dst-port range 200 65535 egress: dst-ip = host-ip-address; dst-port range 200 65535
	ip-prefix/ length	100	ingress: dst-ip = ip-prefix/length; dst-port eq 100 egress: dst-ip = host-ip-address; dst-port eq 100
		200 to 65535	ingress: dst-ip = ip-prefix/length; dst-port range 200 65535 egress: dst-ip = host-ip-address; dst-port range 200 65535

Table 81 **Nas-Filter-Rule and Alc-Nas-Filter-Rule-Shared Fields (Continued)**

Action or Classifier	Value		Corresponding SR OS Filter Function
<options: frag>	frag		fragment true (ipv4 only)
<options: ipoptions>	ssrr		ip-option 9/ip-mask 255
	lsrr		ip-option 3/ ip-mask 255
	rr		ip-option 7/ ip-mask 255
	ts		ip-option 4/ ip-mask 255
	!ssrr		not supported
	!lsrr		not supported
	!rr		not supported
	!ts		not supported
	ssrr,lsrr,rr,ts		not supported
<options: tcpoptions>	mss		not supported
	window		not supported
	sack		not supported
	ts		not supported
	!mss		not supported
	!window		not supported
	!sack		not supported
	!ts		not supported
	mss,window, sack,ts		not supported
<options: established>	established		not supported
<options: setup>	setup		tcp-syn true
			tcp-ack false
			protocol tcp

Table 81 **Nas-Filter-Rule and Alc-Nas-Filter-Rule-Shared Fields (Continued)**

Action or Classifier	Value		Corresponding SR OS Filter Function
<options: tcpflags>	syn		tcp-syn true
	!syn		tcp-syn false
	ack		tcp-ack true
	!ack		tcp-ack false
	fin		not supported
	rst		not supported
	psh		not supported
	urg		not supported
<options: icmptypesv4>	echo reply		protocol 1/icmp-type 0
	destination unreachable		protocol 1/icmp-type 3
	source quench		protocol 1/icmp-type 4
	redirect		protocol 1/icmp-type 5
	echo request		protocol 1/icmp-type 8
	router advertisement		protocol 1/icmp-type 9
	router solicitation		protocol 1/icmp-type 10
	time-to-live exceeded		protocol 1/icmp-type 11
	IP header bad		protocol 1/icmp-type 12
	timestamp request		protocol 1/icmp-type 13
	timestamp reply		protocol 1/icmp-type 14
	information request		protocol 1/icmp-type 15
	information reply		protocol 1/icmp-type 16
	address mask request		protocol 1/icmp-type 17
	address mask reply		protocol 1/icmp-type 18
	-		protocol 1/icmp-type [0 to 255]
	3 to 9 (range)		not supported
	3,5,8,9 (comma separated)		not supported

Table 81 **Nas-Filter-Rule and Alc-Nas-Filter-Rule-Shared Fields (Continued)**

Action or Classifier	Value	Corresponding SR OS Filter Function
<options: icmptypesv6>	destination unreachable	icmp-type 1
	time-to-live exceeded	icmp-type 3
	IP header bad	icmp-type 4
	echo request	icmp-type 128
	echo reply	icmp-type129
	router solicitation	icmp-type 133
	router advertisement	icmp-type 134
	redirect	icmp-type 137

The format for [242] Ascend-Data-Filter and [26-6527-159] Alc-Ascend-Data-Filter-Host-Spec is an octet string with fixed length fields. Table 82 provides details on the respective fields.

Table 82 **Ascend-Data-Filter and Alc-Data-Filter-Host-Spec Fields**

Field	Length	Value
Type	byte	1 = IPv4 3 = IPv6
Filter or forward	1 byte	0 = drop 1 = accept
Indirection	1 byte	0 = egress 1 = ingress
Spare	1 byte	ignored
Source IP address	IPv4 = 4 bytes IPv6 = 16 bytes	IP address of the source interface
Destination IP address	IPv4 = 4 bytes IPv6 = 16 bytes	IP address of the destination interface
Source IP prefix	1 byte	Number of bits in the network portion
Destination IP prefix	1 byte	Number of bits in the network portion
Protocol	1 byte	Protocol number. Match the inner most header only for IPv6

Table 82 **Ascend-Data-Filter and Alc-Data-Filter-Host-Spec Fields**

Field	Length	Value
Established	1 byte	ignored (not implemented)
Source port	2 bytes	Port number of the source port
Destination port	2 bytes	Port number of the destination port
Source port qualifier	1 byte	0 = no compare 1 = less than 2 = equal to 3 = greater than 4 = not equal to (not supported)
destination port qualifier	1 byte	0 = no compare 1 = less than 2 = equal to 3 = greater than 4 = not equal to (not supported)
Reserved	2 bytes	ignored

9.3.14.2 Checking Filter Policy Details

Use following show commands to check filter policy details and the filter configuration for a subscriber host:

```
show filter ip ip-filter-id detail
show filter ipv6 ip-filter-id detail
show filter ip ip-filter-id type entry-type
show filter ipv6 ipv6-filter-id type entry-type
    entry-type : fixed | radius-insert | credit-control-
      insert | radius-shared
show service active-subscribers filter [subscriber sub-
 ident-string] [origin origin]
    sub-ident-string : [32 chars max]
    origin : radius | credit-control"
```

9.3.15 ESM PPPoA/PPPoEoA

The main goal of PPP in the subscriber context is to provide authentication, to negotiate link layer parameters (such as MTU) and to negotiate IP parameters (IP address, WINS, DNS, Default Gateway, and so on).

Each PPP session is carried over a single ATM VC over to the BNG. In PPPoA environment, PPP session is directly encapsulated over ATM transport on a DSL Customer Premise Equipment (CPE).

In the PPPoEoA environment, an additional layer is added to accommodate Ethernet medium. In this fashion, a PPP session can be directly terminated on any host within the customer Ethernet network and then transported over an ATM network to the BNG. Multiple PPPoE sessions can be carried over a single ATM PVC.

However, the majority of current implementations in ATM transport networks have PPPoE session terminated at the DSL CPE and not on a host within the customer network. This PPPoE session is then transported over ATM to the BNG. Although it seems unnecessary to add the overhead associated with Ethernet on an ATM-equipped DSL CPE, mainly for historical reasons PPPoE became ingrained in the home network and as such has moved into DSL CPE.

9.3.15.1 PPPoA

In a PPPoA environment, services are offered to residential and business customers. DSL CPE and BNG are the originating and terminating points of a PPP connection. In a residential example, a DSL CPE is a Home Gateway (HG) as shown in Figure 90.

Figure 90 **PPPoA Architecture and Packet Encapsulation**

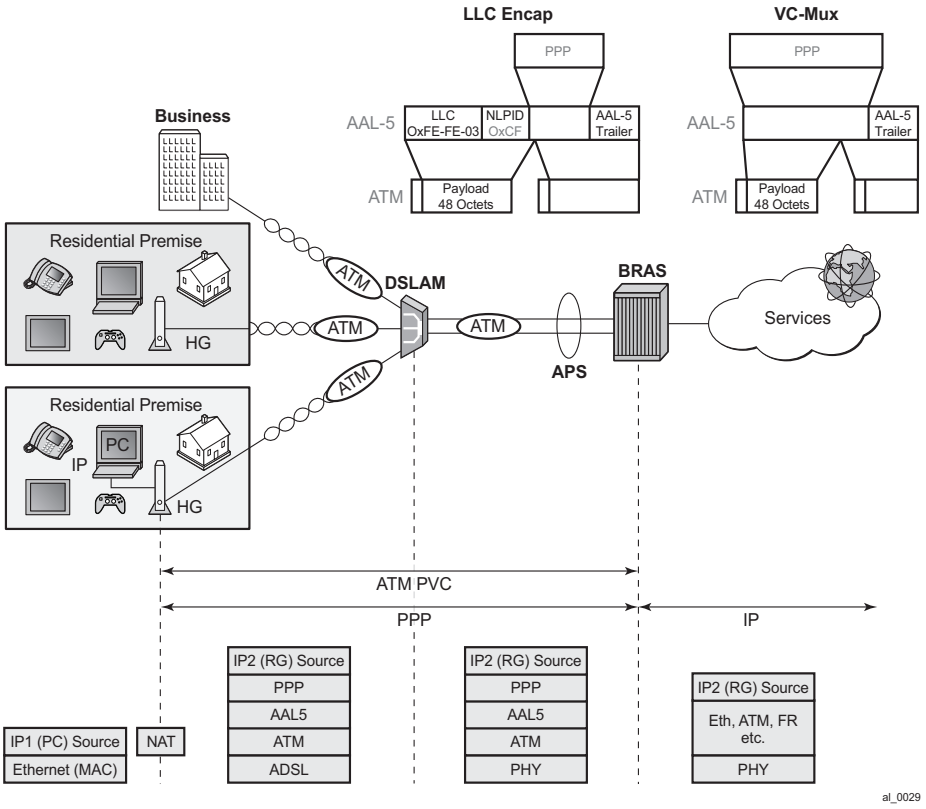

An ATM VC must exist between a CPE and a BNG before a PPP session can be established. From a QoS perspective, this VC is one of the following service categories: CBR, rt-VBR, nrt-VBR or UBR(+). CPE starts negotiating PPP link/session parameters over the VC. Once the IP address is obtained from the Service Provider (SP) side, the customer is ready for data transfer. There is a 1:1 mapping between a PPPoA session and a VC.

There is no need for the customer to run PPP within its own network. CPE, as a default gateway, accepts Ethernet IP packets, strips off the Ethernet header, performs a NAT function, and encapsulates the IP packets into PPPoA before it sends them on to the BNG.

In most residential cases, this PPPoA session is used for pure data transport (Internet access). As such, the BNG side would require a single service queue per VC.

In certain cases, customers use PPPoA for VoIP. In this case, they use an IAD (Integrated Access Device) to gain access to the ATM network.

Nothing in this architecture precludes customers from using multiple services over a single PPPoA session. Services would be differentiated through DSCP bits and each service in this case would require a separate queue. This is a scenario with business customers where a customer runs multiple services over a PPPoA session. DSL CPE would play a role here in differentiating and appropriately marking the services, either through a separate physical port per service or through some other means.

9.3.15.2 PPPoEoA

In the PPPoE model the originating PPP point could extends beyond the DSL CPE and into the customer Ethernet network where any host can originate a PPP session. In such case, ATM is generally used as a transport to carry PPPoE sessions that are originated by hosts within the customer network (beyond DSL CPE). The DSL CPE would operate in a bridged mode. This is shown in Figure 91.

Figure 91 **PPoEoA Host Terminated Session**

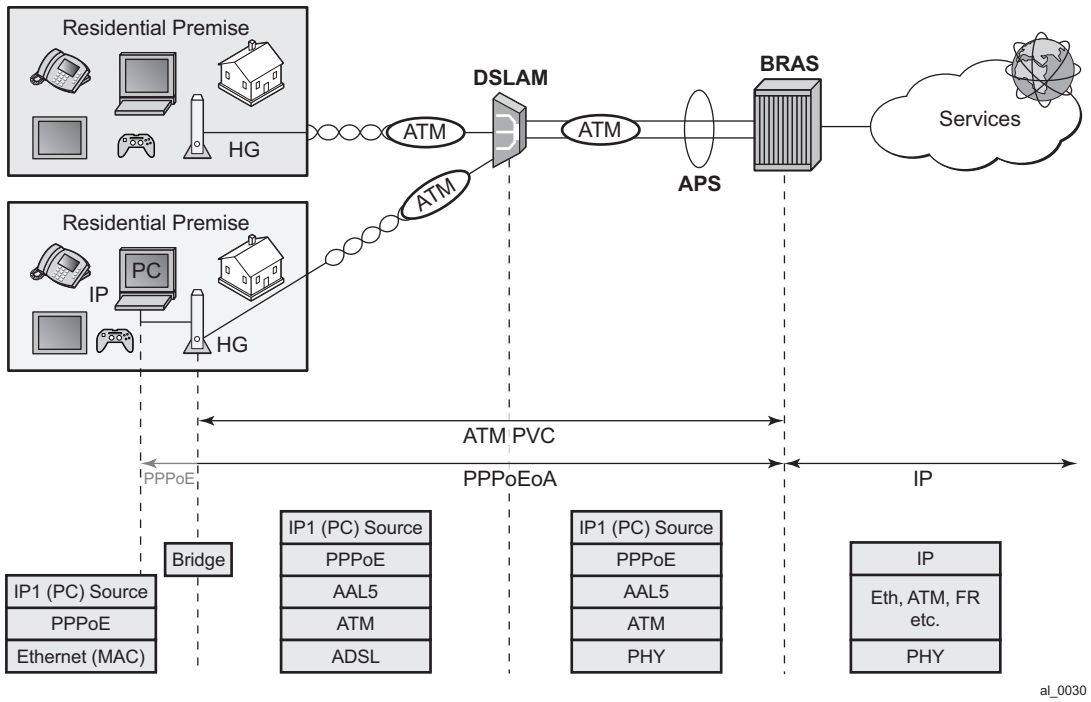

However, the majority of deployments have PPPoE sessions terminated in DSL CPE. Although it is inefficient to add an extra Ethernet encapsulation layer over ATM-equipped DSL CPE, the evolution of PPP in broadband is the main reason for this deployment scenario as shown in Figure 92.

Figure 92 **PPPoEoA DSL CPE Terminated Session**

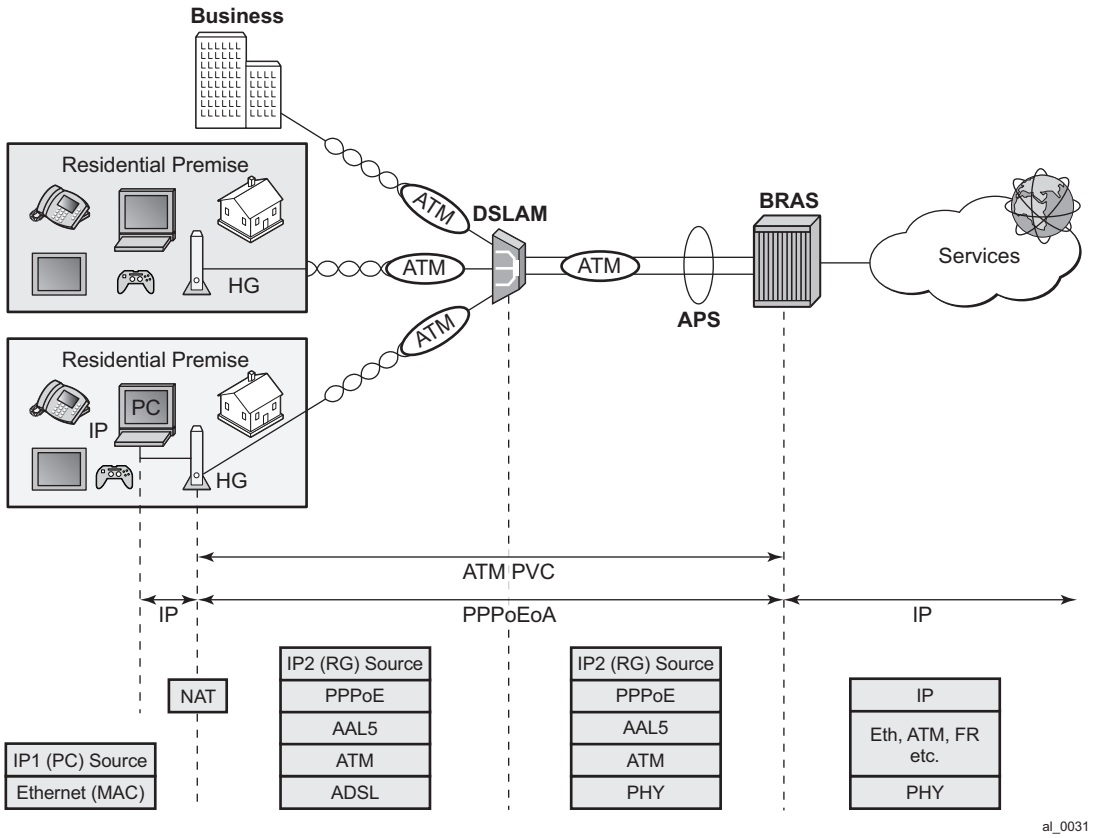

PPPoEoA implementation must allow multiple PPPoE sessions of the same subscriber to be carried over a single ATM PVC.

9.3.15.3 Hardware Support

This feature is supported on ATM MDA on:

- 7750 SR (SR-12, SR-7) platforms
- 7750 SR-c4/12 platforms

ATM MDAs:

- 16 port ATM OC-3/STM-1 (single rate)
- 4 port ATM OC-3/12c/STM-1/4c (dual rate on a per port basis; oc3/stm1 or oc12/stm4; port speed can only be changed in groups of four ports).

PPPoEoA/PPPoA is not supported on the following modules:

ASAP MDAs:

- 4 port channelized OC-3/STM-1
- 1 port channelized OC-12/STM-4
- 12 port channelized DS3/E3 (coax)
- 4 port channelized DS3/E3

ATM CMA:

- 8 port T1/E1 ATM (RJ-48)

The 7750-c4/12 currently supports only the four port ATM MDA.

9.3.15.4 Termination Points

PPPoA/PPPoEoA sessions are terminated on the access ATM SAP in IES and IP-VPN service context through subscriber/group interfaces.

However, for the wholesale/retail deployment scenarios, the ATM VCs are terminated on the LAC while the PPP(oE) sessions will be terminated on the LNS.

PPPoA/PPPoEoA is not supported in the wholesale/retail VRF model (wholesale VRF + retailer VRF). However it is supported in wholesale/retail MSAP model (capture SAP in VPLS that is mapped into a VRF).

9.3.15.5 PPPoA Encapsulation

PPP frames are transported over ATM using ATM AAL5 framing mechanism. This is defined in RFC 2364. In short, each PPP packet is appended by an 8-octet AAL5 trailer with some control information (16-bit length field and a 32-bit CRC being the most important). This new frame is not self-identifying, in other words, kind of payload it carries (PPP, IPv4, IPv6, ARP, MPLS, and so on) can not be identified. This means that if you want to send it as such, it can carry only a single protocol type which must be agreed upon in advance by configuration at each side of a PVC connection. Only then will both ends of the connection be able to recognize the payload type inside of it.

To add more flexibility to the payload type and allow multiple protocol types to be multiplexed within a single VC session, an additional header must be defined in the AAL-5 packet. This header which allows protocol multiplexing over ATM VC is called Link Layer Control (LLC) header. PPP transport over ATM using LLC is defined in RFC 2364, *PPP over AAL5*. A more extensive version for multiplexing protocols over VC is defined in RFC 2684, *Multiprotocol Encapsulation over ATM AAL5" (LLC/ SNAP encap)*.

Two types of PPPoA encapsulation are supported:

1. LLC (RFC 2364)

   LLC header is extended with a NLPID (Network Layer Protocol Identifier) which identifies the protocol type inside of the AAL5 frame. NLPID for PPP in LLC is 0xCF.

Figure 93 **PPPoA LLC Encapsulation**

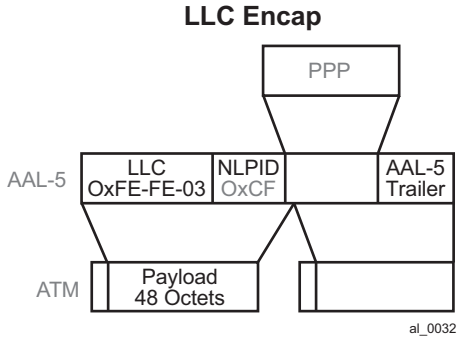

   Possible CLI syntax:

   encapsulation aal5nlpid-ppp

   The **ppp** keyword in the **aal5nlpid-ppp** command is used to indicate that ppp is the only encapsulation that currently supported in NLPID.

   SNAP is an extension to LLC for protocols that are not defined in LLC NLPID. PPP protocol identifier is defined within NLPID and therefore it does not need additional SNAP header (this reduces the overall AAL5 header overhead).

   LLC/SNAP encapsulation is defined in RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, and it is needed for PPPoEoA encapsulation.

2. VC-MUX

   In VC-MUX mode there is no additional (LLC/SNAP) header used for protocol multiplexing. Instead, VC endpoints must agree beforehand on the payload type that they will transport. For PVCs this is done during the provisioning phase on each side of the connection. For example:

   encapsulation aal5mux-ppp

aal5mux-ppp will tell each end of the VC in advance that the payload inside of the AAL5 frame is PPP.

Figure 94 **PPPoA AAL5MUX Encapsulation**

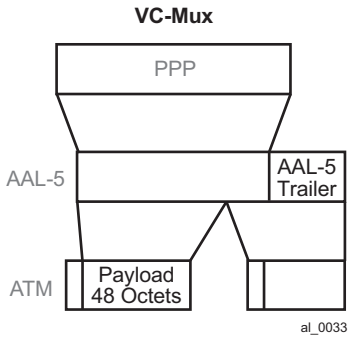

9.3.15.5.1 PPPoEoA Encapsulation

Similar to PPPoA, two types of encapsulation are defined for PPPoEoA:

- Protocol multiplexing (Layer 2 bridging defined by additional headers - LLC/SNAP).
- VC-Multiplexing where payload type within a VC is agreed upon in advance by static configuration.

Both types are supported in our implementation:

1. Multiprotocol Encapsulation over ATM AAL5 (LLC/SNAP) defined in RFC 2684.

   This encapsulation adds support for protocols that are currently not defined in the LLC header. There are two basic types of encapsulations defined under this RFC:

   – Routed

   – Bridged

   PPPoEoA encapsulation is defined as Bridged where Layer 2 information is preserved through transitioning between two different Layer 2 network types (Ethernet -> ATM) which is shown in Figure 95.

   0xAA-AA-03 in the LLC header indicates the presence of the SNAP header.

   0x00-80-C2 in the OUI indicates that a bridged PDU is encapsulated.

   0x00-01 or 0x00-07 in PID indicates that the encapsulated Layer 2 network type is 802.3 Ethernet with or without preserved FCS.

Figure 95 **PPPoEoA Bridged LLC/SNAP Encapsulation**

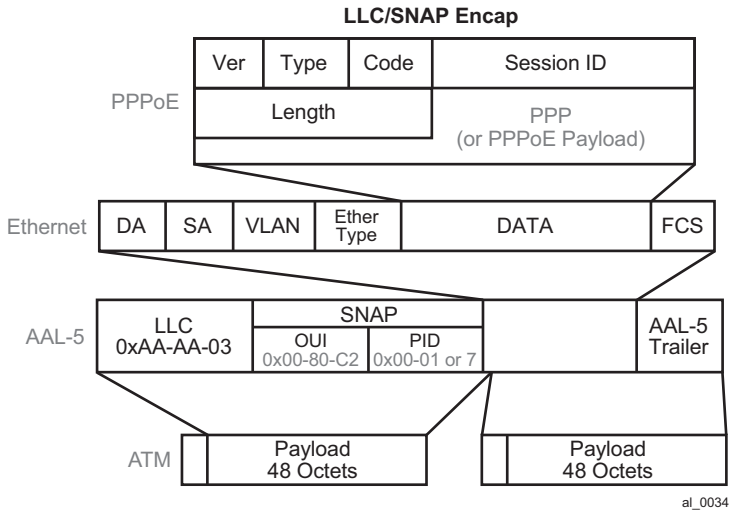

The CLI syntax is:

encapsulation aal5snap-bridged

- VC-MUX

  In the VC-MUX mode there is no additional (LLC/SNAP) header used for protocol multiplexing. Instead, the VC endpoints must agree before hand on the payload type that they will transport. For PVCs this is done during the provisioning phase on each side of the connection. For example:

  encapsulation aal5mux-bridged-eth-nofcs

  aal5mux-bridged-eth-nofcs tells each end of the VC in advance that the payload inside of the AAL5 frame is an Ethernet frame. In this case, it accepts the frame and treat it as an Ethernet frame inside AAL5. The EtherType within the frame must be set to 0x8863 (PPPoE Discovery Phase) or 0x8864 (PPPoE Session Phase).

  The '-nofcs' portion indicates that the FCS is not supported on those Ethernet frames.

  PPPoEoA encapsulation is shown in Figure 96.

Figure 96 **PPPoA AAL5MUX Encapsulation**

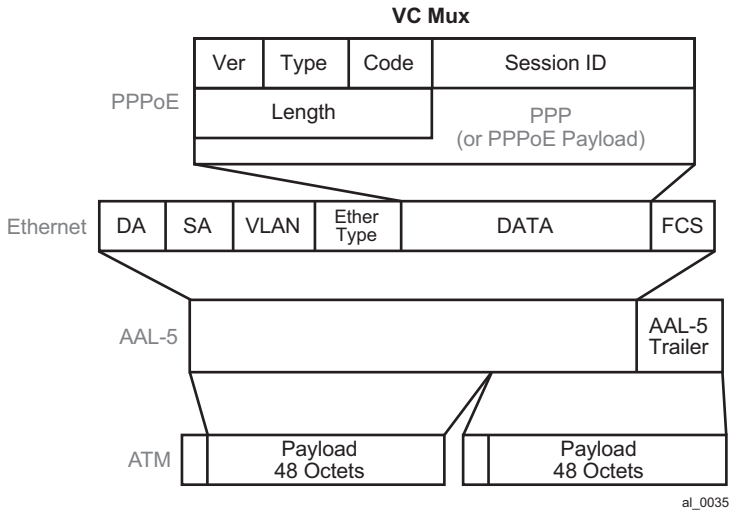

9.3.15.6 Encapsulation Summary

Figure 97 **LLC/SNAP Encapsulation**

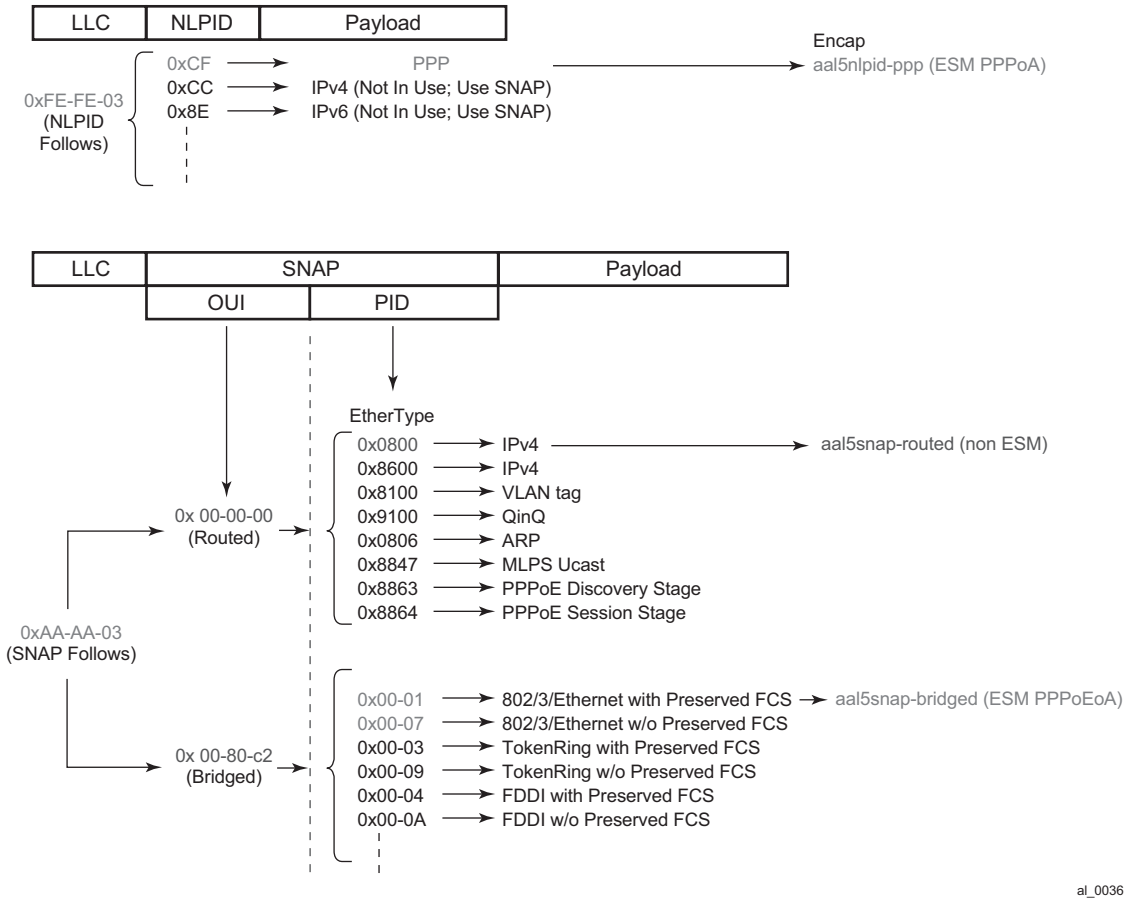

There is a 0-2 byte additional padding (optional) in the SNAP Bridged encapsulation header that is not shown in Figure 97. This padding, according to RFC 2684, is necessary to align the info field (payload of the Layer 2 encapsulated frame) on a 4-Byte boundary.

At the end the following types of ATM encapsulation are supported in ESM:

- aal5snap-bridged —This encapsulation type is used for PPPoEoA encapsulation with FCS or without FCS on ingress. On egress only frames without MAC FCS can be sent. PPPoE session type will be determined based on the EtherType in the Ethernet frame.
- aal5mux-bridged-eth-nofcs — This encapsulation type is used for mux PPPoEoA sessions without MAC FCS.

- aal5nlpid-ppp — This encapsulation type is used for LLC/NLPID PPPoA encapsulated packets.
- aal5mux-ppp — This type encapsulation is used for PPPoA traffic without LLC/SNAP header (VC-MUX).
- aal5auto — This encapsulation type is supported in ESM and it is used for auto detecting the encapsulation type. This is also called autosensing.

All hosts for the same subscriber use the same encapsulation type.

9.3.15.7 Concurrent Support for Different Service Types on the Same Port

An ATM port on 7450 ESS and 7750 SR can concurrently support various service types. One service type can be mapped only to one VC. This would normally be the case if there is an aggregation network in front of 7450 ESS and 7750 SR. A service provider could run a variety of services over this aggregation network (ATM switches) connected on the same physical port within the routers.

For example, the following services can be run on the same physical port:

- PPPoA or PPPoEoA sessions connecting residential or business customers through a DSLAM to 7750 SR and 7450 ESS
- Providing access for Xpipe services
- Plain aal5snap-ip or vc-mux-ip PVC for Internet access

9.3.15.8 Restrictions in Scaled ATM MDA Mode

ATM concatenation mode for Apipe is not supported in the 16K VC mode.

In the concatenated mode, cells are delayed so that they can be concatenated and delivered in a single packet over to pseudowires (PWs) to the other side. Without concatenation, each cell is transported individually. The implication is that each cell is individually encapsulated into Eth/MPLS which results in wasted bandwidth on the link.

Support for the concatenated ATM pseudowires is not removed from the CLI in the 16K mode of operation:

```
configure
   service apipe <id> [vc-type <cell-type>]
     spoke-sdp <sdp:pw> cell-concatenation
        [no] aal5-frame-aware
```

```
[no] clp-change
[no] max-cells
[no] max-delay
```

Instead:

- Adding an ATM port or a connection profile on an MDA in 16-VC mode to an APIPE is disabled if the vc-type is set to atm-cell and cell-concatenation is enabled.

  A:BNG>config>service>apipe# sap x/y/z:cp.w create

  MINOR: SVCMGR #2603 Cell-concatenation is not allowed on 16k VC-mode ATM MDAs
- cell-concatenation on an APIPE of the vc-type atm-cell is disabled if it already contains an ATM port or a connection profile.

  A:BNG>config>service>apipe>spoke-sdp>cell-concat# max-delay X

  MINOR: SVCMGR #2603 Cell-concatenation is not allowed on 16k VC-mode ATM MDAs

Note that regular VPI/VCI SAPs (sap:vpi/vci) are not allowed to be configured on an Apipe of vc-type atm-cell.

Cell-concatenation is supported on Apipe services with a VC on an ATM MDA in the 8-VC mode.

AAL5 SDU mode is continued to be supported.

9.3.15.9 QoS Implementation

In addition to our system QoS that is provided in the “Q” chip on IOM, the ATM MDA offers QoS capabilities at the ATM cell level. In the context of this document, the system QoS is referred to as ‘IOM QoS’ and to ATM MDA provided QoS simply as ‘ATM QoS’.

Both, IOM QoS as well as ATM QoS defined by Traffic Descriptors working at the cell level, play a role in the overall QoS for the SAP (or virtual circuit). ATM QoS defines rates of each VC stream and defines the behavior under port congestion. IOM QoS defines bandwidth allotment and the scheduling scheme for each service within a VC stream.

In general, both MDAs, ATM and ASAP, support four traffic categories:

- CBR
- Rt-VBR

- Nrt-VBR
- UBR

Figure 98 **Scheduling on ATM MDA**

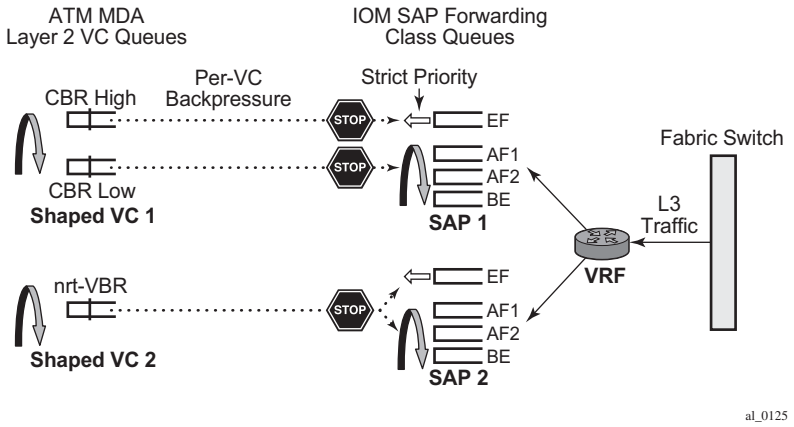

Each ATM traffic category is defined by a set of parameters, such as PIR, SIR, MIR, MBS and CDVT.

Currently, policing is supported only on ingress for CBR and VBR traffic classes. CBR traffic class police at PIR, while rt/nrt-VBR police at SIR.

Shaping is supported only on egress for CBR and rt/nrt-VBR traffic classes. CBR shapes traffic at PIR, while rt/nrt-VBR shape traffic at SIR. Egress shaping can be disabled only for nrt-VBR traffic class.

Scheduling at the ATM layer is shown in Figure 98. Shaped CBR and VBR traffic classes have two queues, an HP and an LP queue. Packets from IOM are marked according to scheduling priority of the Forwarding Class (expedited or best-effort) from which they were sent and are accepted into the ATM VC queue (HP or LP) accordingly. For example, at the IOM level packets from an expedited FC (queue) are marked as HP, and the packets from a best-effort FC (queue) are marked as LP. When these packets arrive to MDA, they are admitted into appropriate queues, pending VC (or port) buffer availability.

Non-shaped VBR and UBR have only one queue at the MDA level.

In terms of the ATM QoS scheduling, CBR has the highest scheduling priority (strict priority) followed by rt-VBR (also strict priority). The remaining traffic classes (non-shaped VBR and UBR) are serviced in WRR fashion, the weight being their configured SIR (non shaped nrt-VBR) or MIR (UBR) rate.

There are several areas of QoS that are addressed in relation to integration of IOM QoS and ATM QoS:

- Association between the subscriber and ATM VC traffic descriptor.
- Rate adjustments between Layer 3 and ATM QoS rates due to difference in frame/cell sizes on which they operate.
- Per VP shaping.

9.3.15.10 Association Between the Subscriber and ATM VC Traffic Descriptor (QoS)

Each PPPoA PVC is of a certain traffic class – CBR, rt/nrt-VBR or UBR/UBR+MIR which along with other parameters is defined in the ATM traffic descriptor profile. For ATM service categories and traffic descriptors. There should be a uniform mapping between service offered and Layer 3 and ATM QoS.

Figure 99 **ATM Traffic Descriptor Association with Subscriber**

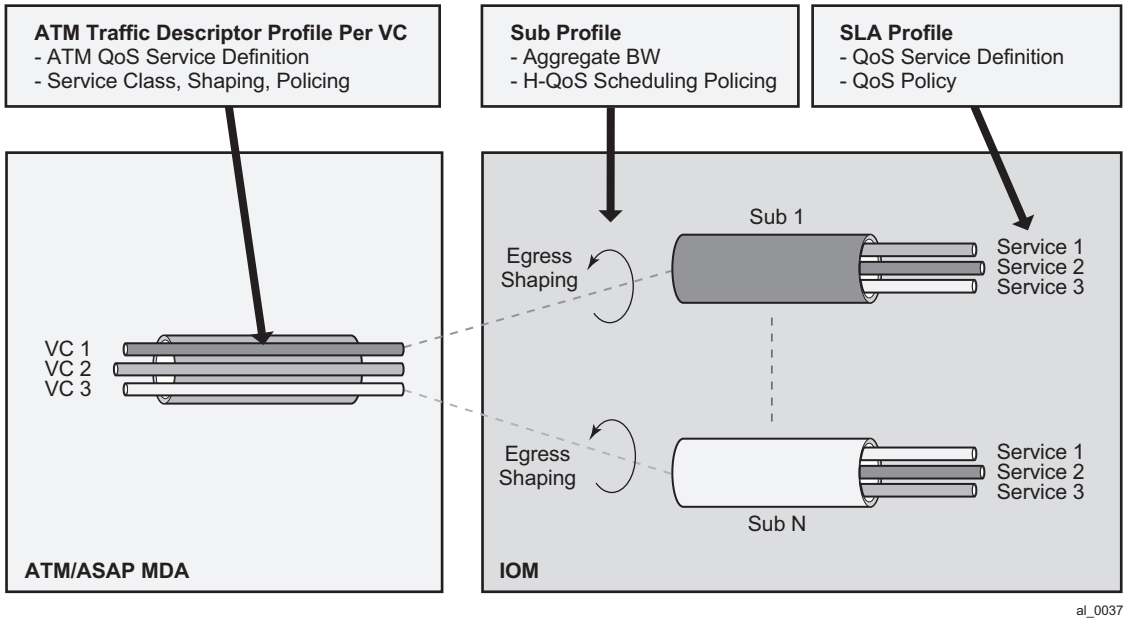

The ATM traffic descriptor (atm-td) is applied to a VC under the SAP CLI hierarchy. In this fashion, MDA related QoS (ATM QoS) is referenced outside of the subscriber context (SUB/SLA-profiles).

The following describes the operability:

atm-td for the VC is applied under the SAP:

```
configure
    services ies/vprn
        subscriber-interface <sub-if-name>
            group-interface <grp-if-name>
                   sap <sap-id>
                       atm
                           ingress
                               traffic-desc <id>
                           egress
                               traffic-desc <id>
```

and for MSAP:

```
configure
    subscriber-management
        msap-policy <name>
            atm
                   ingress
                       traffic-desc <id>
                   egress
                       traffic-desc <id>
```

The msap-policy is then invoked via LUDB, RADIUS or the default-msap-policy under the capture SAP.

This allows router to have preconfigured MSAP policies, each corresponding to a specific VC type with its own traffic-class parameters (CBR/rt-nrt-VBR/UBR.). Each subscriber with corresponding hosts is then associated with a msap-policy that determines the VC type.

A more flexible way to go about this would be to allow the subscriber to reference the atm-td directly by the sub-host at the host creation time, independently of the msap-policy. This is supported in the following manner:

- atm-td is referenced via a RADIUS VSA in the Access-Accept message.
- There are two VSAs, one for ingress and one for egress atm-td:
  - alc-ingress-atm-td
    (ATTRIBUTE Alc-ATM-Ingress-TD-Profile 128 integer)
  - alc-egress-atm-td
    (ATTRIBUTE Alc-ATM-Egress-TD-Profile 129 integer)
- Only the first host of the subscriber can overwrite the atm-td that is defined under the msap-policy or under the static SAP.

- Consecutive hosts (second, third, and so on) of the same subscriber will not have any effect on the atm-td of the VC. For example, if the second host tries to overwrite the existing atm-td with a different atm-td, it will fail. Once the atm-td is set via the Access-Accept message for the first host, it cannot be changed as long as the subscriber is active in the system. This implies that all hosts of the same subscriber will have the same atm-td.
- If another hosts of the same subscriber initiates a session with a different atm-td name in the Access-Accept message, the host will be created but the atm-td for the VC will not be changed. A trap/syslog will be generated with a 'VC-parameter-mismatch:VSA ignored' info message.
- atm-td change will not be supported through CoA.

The following guidelines should be followed in configuring ATM traffic descriptors:

- ATM policing is the only function that can discard ATM cells on an ATM MDA
- Ingress:
  - ATM policing can only be enabled on ingress.
  - Policing is disabled by default.
  - Policing can only be enabled for the following traffic categories: CBT, rt-VBR and nrt-VBR.
  - CBR police at PIR. CBR is CLP transparent (it does not look at the CLP bit – aggregate traffic policing).
  - rt-VBR and nrt-VBR police at SIR and can operate either in CLP Significant (CLP marking) or CLP transparent mode (only relevant for Apipes and not for ESM).
  - In CLP significant mode with policing enabled, traffic is policed at the PIR value, but it can be marked with CLP 1 if it exceeds the SIR rate (only relevant for Apipes and not for ESM).
  - Ingress classification: the service category and CLP bits only affect Apipe traffic. If ATM traffic is terminated in the BNG, rely on the existing QoS classification implemented in the IOM.
- Egress:
  - Shaping is supported only on egress. Traffic shaping is supported for CBR, rt/nrt-VBR service categories.
  - ATM MDA is cell "lossless" — It sends backpressure to IOM which drops packets.

9.3.15.11 Per VP Shaping

VCs traversing the same DSLAMs typically uses the same VPI. To prevent overrunning DSLAM capacity (intermediate destinations) in case an aggregation network is in place (between BNG and a DSLAM), per VP shaping will be implemented.

Figure 100 **VP Shaper**

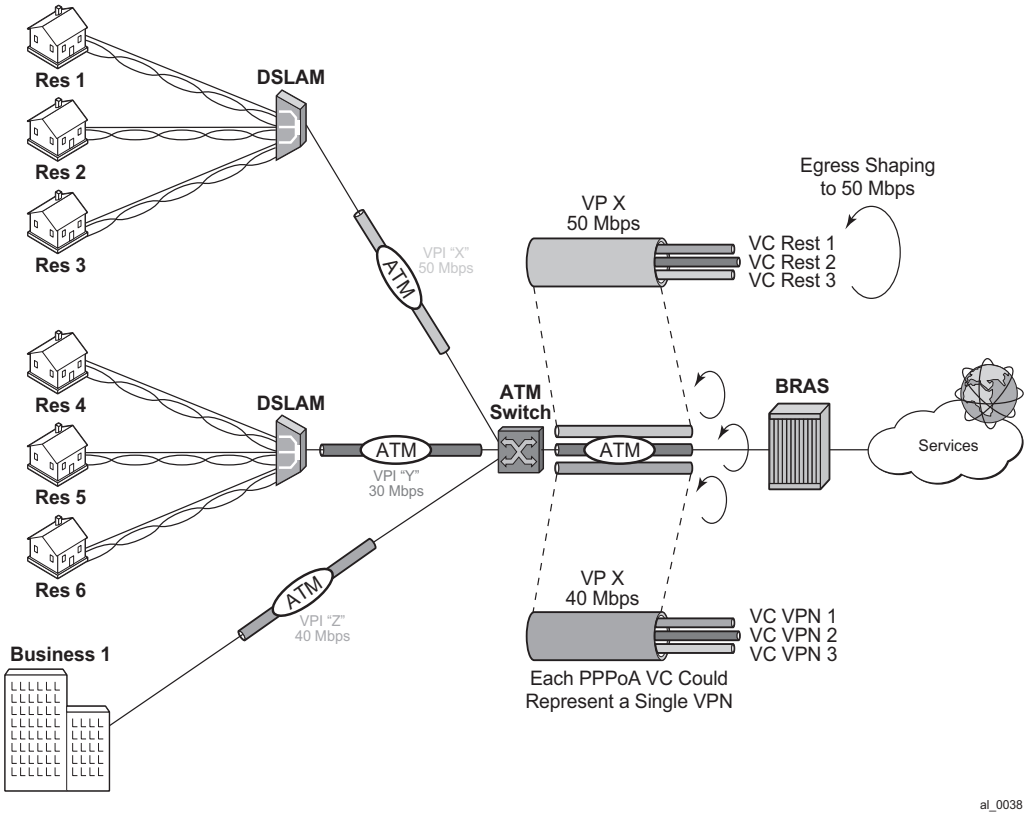

The shaping rate per VP and the VP service type is provisioned manually via a traffic descriptor on a per port level. The CLI syntax is:

```
config
    port <port-id>
        sonet-sdh
            path
              atm
                vp <vpi> egress-traffic-desc <atm-td-profile-id>
```

Where the vpi is a VPI identifier and the egress-traffic-desc is the traffic descriptor id. Only traffic descriptors with service-category of cbr, rt-vbr and nrt-vbr can be used in VP Shapers. However, only the rate in the traffic descriptor can be changed on-the-fly, and not the service category (nrtVBR, rtVBR, CBR). A VP shaper can be added to or removed from active VCs.

CBR VP Shaper shapes cells at the exact PCR rate. There is no burst concept in CBR shaping. Excessive traffic is back-pressured towards the IOM (Q-chip).

The IOM will never send more frames to the ATM MDA than the MDA cannot buffer. This is implemented through a combination of software and hardware backpressure mechanisms. This backpressure mechanism utilizes a combination of hardware and software. Software backpressure aims to have around 100ms of traffic queued against a VC based on its configured shaping/scheduling rate, but being a software mechanism that is only a guideline. As ATM MDA detects more traffic that it can accept, a hardware backpressure is exerted.

A rt/nrt-VBR type VP Shaper has three parameters associated with it: PCR (peak cell rate), SCR (sustained cell rate) and MBS (maximum burst size at a peak rate). As long as there is enough MBR credit, traffic will be shaped at the PCR rate. Once all MBR credit (burst) is exhausted, traffic will be shaped at the SCR rate. Bursting above the SCR is configurable via the MBS parameter.

In both cases cells will be spaced at 1/PCR or 1/SCR as perfectly as possible with minimum jitter.

VP shaping is supported only when the ATM MDA is in max16k-vc mode. The maximum number of VP shapers per MDA to 128.

The maximum number of VCs that can feed into a single VP shaper is 16K. This includes the sum of all VC-ranges on the VP plus any statically configured VC on that VP.

VCs within the VP tunnel is serviced by a single scheduler assigned to each VP tunnel. The ATM VP shaper will condition the aggregate traffic for all ATM VCs within the VP tunnel. VCs within the shaped VP tunnel are degraded from the originally assigned service category to a common UBR service category (default traffic descriptor). If the VP shaper is removed from the VCs, the VCs will be reverted to their original service category. Scheduling between VCs will be WRR based with a weight parameter that is explicitly configured. The weights assigned to VCs within the VP tunnel are in range 1-255. By default, VCs are assigned a priority based on the originally assigned service category:

- VC degraded from CBR = weight 10
- VC degraded from rt-VBR = weight 7
- VC degraded from nrt-VBR = weight 5

- VC degraded from UBR+ = weight 2
- VC degraded from UBR = weight 1

The weight parameter is user configurable under the traffic descriptor hierarchy.

```
configure
    qos
        atm-td-profile <td-profile-id> [create]
            weight <weight>

        <weight> : 1-255.
```

If weight is not specifically configured, the defaults are taken as described above.

The explicitly configured weight parameter is honored only on ATM MDA in the max16k-vc mode. On all other ATM capable MDAs (ASAP or ATM MDA in max8K-VC mode), the weight parameter is ignored.

Note that in the current ATM implementation there is already a WRR scheme in place based on internally calculated weights. This WRR scheme is used to service traffic from the VC queues of equal priority (where there are two queues per VC - a HiPrio and a LowPrio queue). Weights are assigned to VCs automatically based on the rate of the VC,

9.3.15.12 ATM/IOM QoS Integration

There are major differences between the QoS mode of operation at the ATM MDA level and the IOM level.

ATM QoS operates on fixed size cells that contain additional transport overhead. In addition, ATM shaping is very accurate so that traffic is paced into the ATM network with nodes that are sensitive to bursts (ingress policing). Buffering and service differentiation (number of queues) at the ATM layer is not as flexible as it is on the IOM level.

On the other hand, HQoS in the IOM is less accurate and less responsive to sudden traffic fluctuations (bursts). It operates on Layer 2 frame lengths. Bursts of traffic are usually let into the network more freely than the ATM network would like to accept.

To combine the extensibility of IOM HQoS with the stringent ATM QoS requirements, the two modes of operation are integrated.

The congestion in the ATM network is treated by using extensive IOM HQoS.

When the VC ATM queue becomes congested, it exerts backpressure to the subscriber queues in the IOM on the corresponding VC.

To avoid the condition where a VC becomes overly oversubscribed and excessive in exerting the backpressure, our HQoS in IOM has to be proactive and has to be able to detect congestion on the IOM level based on the ATM VC configured rate or the VP configured rate. IOM HQoS must deal with this congestion before it actually becomes the problem in the ATM layer. Occasional and short-lived backpressure from VCs, but persistent QoS backpressure and congestion at the ATM layer is avoided. A prerequisite for treating ATM congestion at the IOM level is to adjust the frame size in HQoS calculations (on all levels - including queues and aggregate rate limits) so that it reflects the ATM overhead associated with ATM transport (AAL5 encap, cellification).

9.3.15.12.1 Intermediate Node Rate Limit/Shaper

In many cases it is desired or even necessary to shape traffic per DSLAM. In the ATM network, a DSLAM corresponds to a VP (Virtual Path) which is at the ATM cell level identified by a Virtual Path Identifier (VPI). The VPI/VCI pair represents the addressing mechanism in the ATM network.

Another level of hierarchy is introduced into our HQoS model - an aggregate rate limit in the IOM that will correspond to the ATM VP shaper. The purpose of this new construct in the IOM is to help detect congestion at the IOM level before it becomes a severe problem at the ATM layer.

As a final solution, the HQoS hierarchy in the IOM is a 4-tier hierarchy consisting of:

- subscriber queues that are parented to the port scheduler
- egress agg-rate-limit that represents the egress aggregate subscriber rate
- tier-2 aggregate-rate-limit that represents the ATM VP rate limiter (related to DSLAM). This is a pure rate limiter and not a scheduler (it does not receive any scheduling opportunity).
- port-scheduler to which are subscriber queues parented. The port-scheduler delegates bandwidth to its children based on the priority levels associated with the children queues.

The four tiered hierarchy is needed to deal with:

- port congestion. The port bandwidth can be overbooked the desired service levels can still be ensured.
- ATM VP congestion. An ATM VP can be effectively protected from prolonged congestion that would result into significant backpressure to the IOM queues.
- subscriber congestion. Within the subscriber, bandwidth is managed within the subscriber bandwidth limit.

The four-tiered hierarchy looks like Figure 101.

Figure 101 **Tier HQoS**

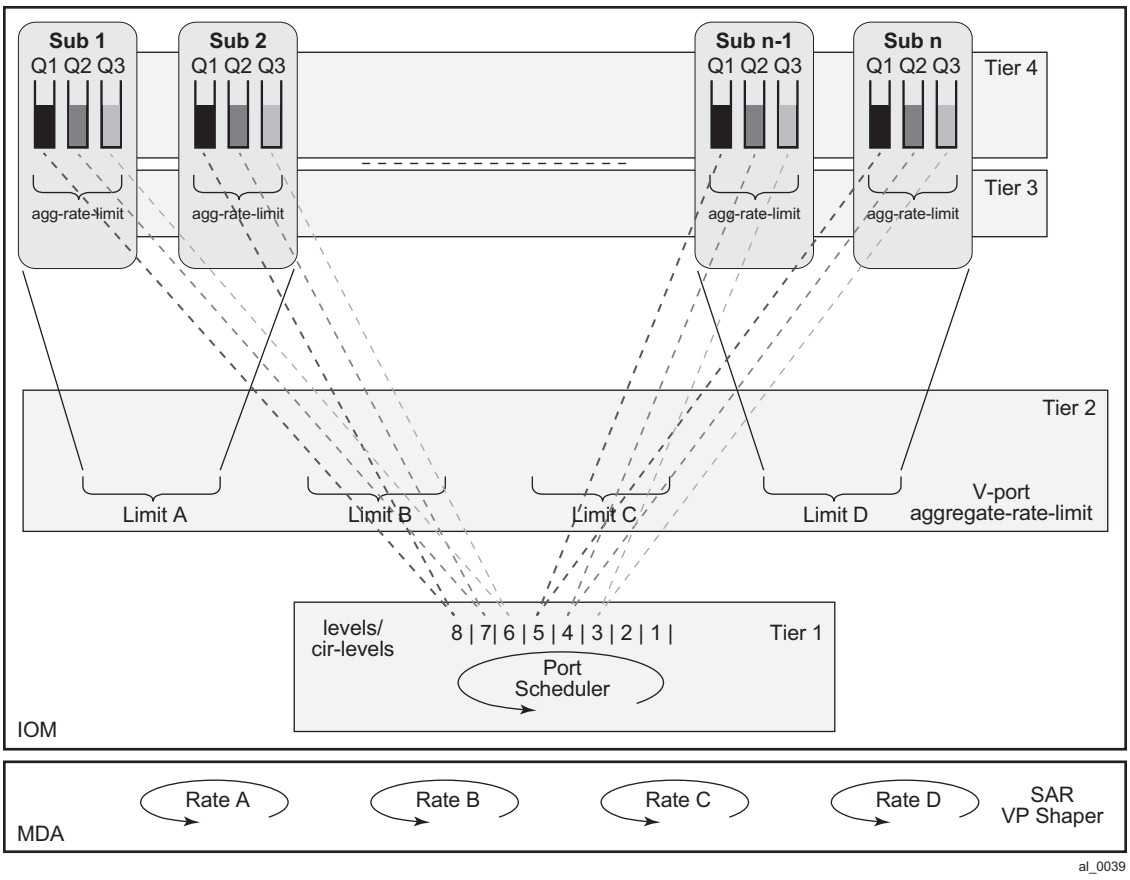

The key point in such HQoS model is that the port-scheduler delegates its available bandwidth to the subscriber queues directly according to the queue priority on a configured level. Higher priority queue are served over all subscribers before any lower priority queues, up to the limits imposed by the tier 1 and tier 2 aggregate rate limits.

9.3.15.12.2 Provisioning Aspects

The tier 2 aggregate rate limit that corresponds to the VP shaper on the IOM level is provisioned in the context of a Vport. The Vport is a container that is configured directly under the port and it can contain either an aggregate rate limit or another port scheduling policy (port scheduler). These two configurations, **vport aggregate** and **vport port scheduling** policy) are mutually exclusive. In addition, if a Vport port scheduling policy is configured instead of the Vport aggregate, then the port scheduler cannot be used (currently two port schedulers applied under the same port cannot be used in a parent/children relationship).

The V-port aggregate have to have the **agg-rate-limit** defined explicitly. In other words, the rate **cannot** be implicitly inherited by configuration from the VP shaper traffic descriptor.

Example:

```
config
    port <port-id>
        sonet-sdh
            path
                egress-scheduler-policy <port-scheduler-policy-name>
                access
                    egress
                        vport <vport-name>
                            description <description-string>
                            host-match dest <dest-string>
                            agg-rate-limit <agg-rate>
                            port-scheduler-policy <port-scheduler-policy-name>

                atm
                    vp <vpi> egress-traffic-desc <atm-td-profile-id>
```

The association between the Vport aggregate and the subscriber host is done in three ways:

- RADIUS attribute — a dest-string VSA that is matched against the string defined under the corresponding Vport. This dest string VSA is returned via RADIUS at the subhost instantiation time.
- LUDB – similar to the RADIUS method. The dest-string comes from the LUDB instead from RADIUS.
- VP identifier (VPI) - The VPI is known to the CPM from the beginning of the session initiation process – raw (unknown encapsulation) packets and passed to the CPM. These raw packets contain VPI,VCI identifiers. The Vport container for the subscriber is referenced implicitly via the VPI in the following fashion (Figure 102):
  - CPM determines the VPI from the first packet for the session.

– ATM VP Shaper mapping: The VPI is used to make the subscriber association with the VP Shaper which is defined under the VP Shaping node: **port>sonet-sdh>path>atm>vp**. The VP Shaping node name must be the VPI number.

– Vport mapping: The *vport-name* in the **port>sonet-sdh>path>access>egress>vport** *vport-name* hierarchy is matched against the VPI number.

Figure 102 **VPI Based V-Port <-> Subscriber Association**

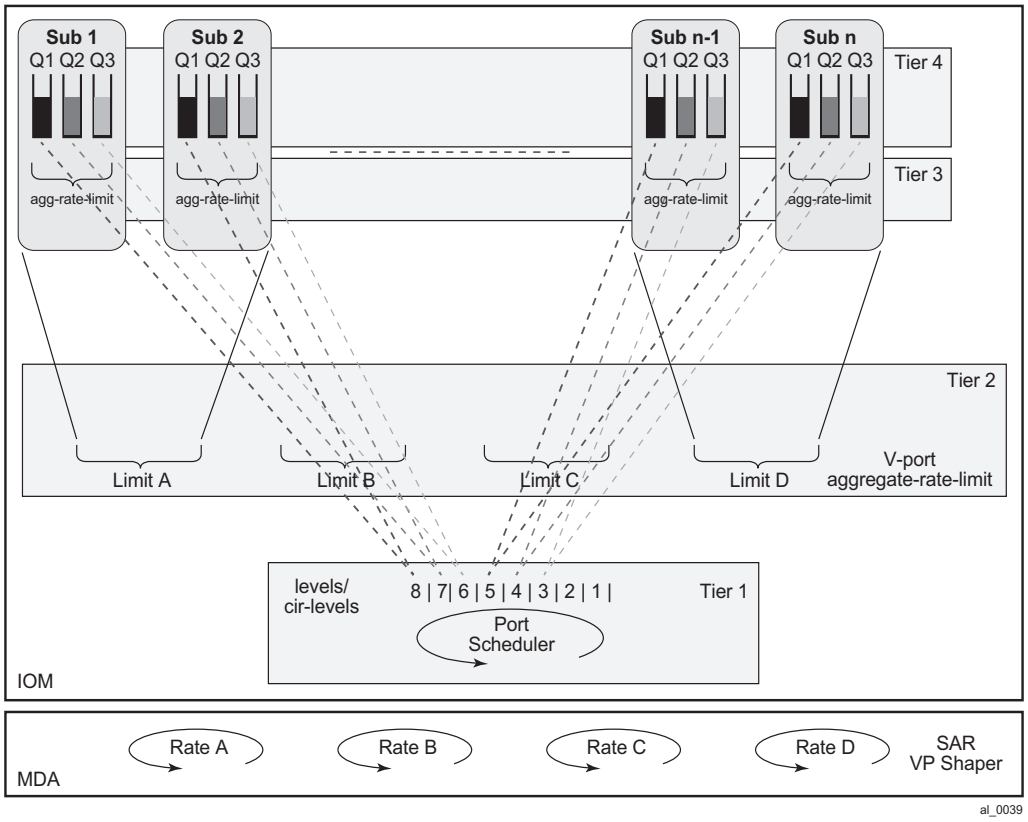

The association method (automatic via VPI or based on RADIUS/LUDB) between the subscriber host and the Vport is defined under the SAP (or MSAP) where the subscriber resides. In most cases subscriber management is used with MSAPs.

This is the syntax:

```
configure
    subscriber-mgmt
        msap-policy <name>
            sub-sla-mgmt
        def-inter-dest-id string <inter-dest-string>
        def-inter-dest-id {use-top-q  | use-vpi}
```

```
configure
    services ies/vprn
        subscriber-interface <sub-if-name>
            group-interface <grp-if-name>
                sap <sap-id>
                    sub-sla-mgmt
                        def-inter-dest-id string <inter-dest-string>
                        def-inter-dest-id {use-top-q | use-vpi}
```

The **def-inter-dest-id** stand for a ‘default inter-destination identifier’.

If the use-vpi method is used and the VPI derived from the incoming traffic points to a non-existing VP container, the association between the subhost and the V-port container will fail and a message/trap is logged. This however will not prevent the creation of the subhost that can be parented to the port-scheduler.

If the use-vpi is used on an Ethernet port where this parameter is not applicable, the parameter is ignored and defaulted to ‘string’. A message/trap is logged.

Note that if the Vport contains an aggregate-rate-limit, then there is no need for the indication of a Vport configuration in the sub-profile or sla-profile of the subscriber. If that the Vport contains a port-scheduling-policy, the sla-profile template must contain the indication that the subscriber is tied to a Vport (**config>subscr-mgmt>sla-profile>egress>qos sap-egress-qos-policy-id vport-scheduler**.)

9.3.15.12.3 HQoS Combinations

These are the possible HQoS combinations that are supported:

- port-scheduler assigned to path, agg-rate-limit assigned to vport, agg-rate-limit assigned to subscriber.
- port-scheduler assigned to path, no vport, agg-rate-limit or scheduler-policy assigned to subscriber.
- no port-scheduler assigned to path, port-scheduler assigned to V-port, agg-rate-limit or scheduler-policy assigned to subscriber.
- no port-scheduler assigned to path, no vport, scheduler-policy assigned to subscriber.

In case that the Vport contains the agg-rate-limit, any subscriber host queue that is parented to a virtual scheduler will not be rate-limited by the Vport aggregate rate. The queue will compete for bandwidth directly on the port's port scheduler, at the priority level and weighted scheduler group the virtual scheduler is port-parented to. If the virtual scheduler is not port-parented or if there is no port scheduler policy on the port, the host queue will be orphaned and will compete for bandwidth directly based on its own PIR and CIR parameters.

9.3.15.12.4 ATM Rate Adjustment

The difference in cell/frame overhead on the ATM level and IOM Level (Layer 2) lead to inconsistent behavior in integrated QoS. For example, IOM based QoS operates on Layer 2 frames (PPP header is included in rate calculations). On the other hand, ATM QoS operates on ATM 53byte cells. Each PPP frame with AAL5 overhead is segmented into 48-byte chunk cells prepended by a 5 byte ATM header. Assuming that ATM QoS operates on 53-octet cells, a significant rate discrepancy might arise from this difference.

For example, consider two flows:

Flow 1, 1000pps, PPP+IP packet size 190Bytes => L2 (IOM) QoS rate is 1520 kb/s

Flow 2, 1000pps, PPP+IP Packet size 232Bytes => L2 (IOM) QoS rate is 1856 kb/s

Assuming that VC MUX encap is used, and that AAL5 trailer (8 octets + padding) is used:

Flow 1 = AAL5 length is 190+8+48B_boundary_padding= 240Bytes => 5 ATM cells

Flow 2 = AAL5 length is 232+8+48B_boundary_padding= 240Bytes=>5 ATM cells

Each cell has a 5 byte header which translates into a 2120 kb/s rate for each flow.

As a result, more bandwidth is required at the ATM layer than the QoS on the IOM level has budgeted. Moreover, the two flows appear to be of the same rate at the ATM layer, even though the IP rate of flow 1 is ~20% less than the flow 2.

From the example above, Flow 1 consumes 40% more bandwidth configured at the ATM layer than it is given in IOM QoS. The rate adjustment depends on the packet size which additionally complicates conversion. This issue is more pronounced with smaller packet sizes.

An example is shown in Figure 103.

Figure 103 **QoS Adjustment**

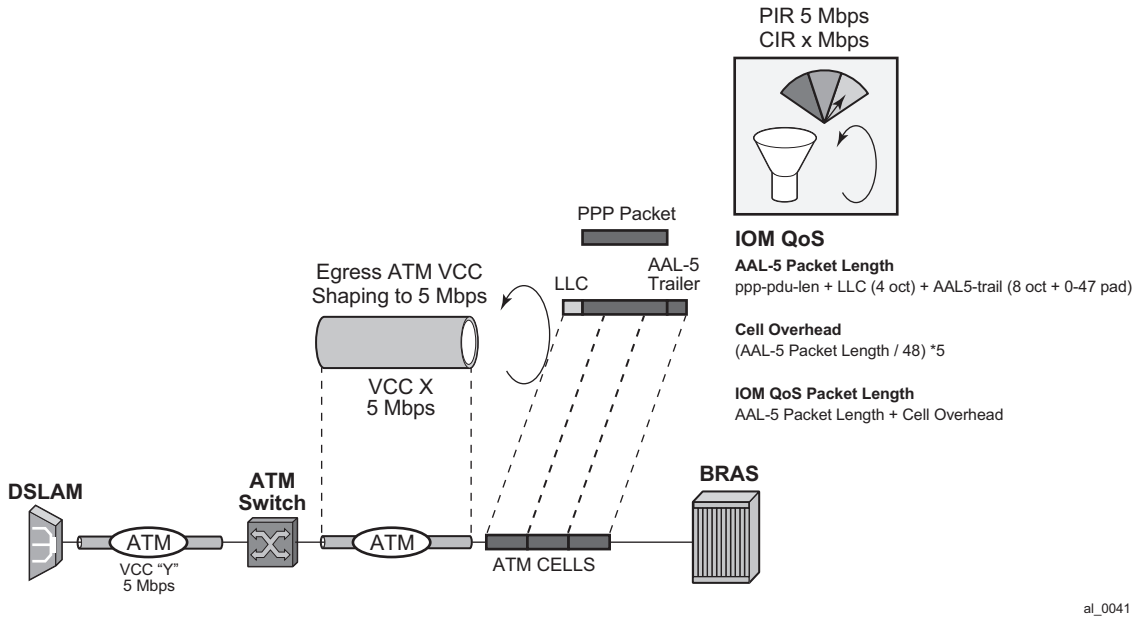

Depending on the session encapsulation (VC-MUX, LLC/SNAP) the packet length on which IOM QoS operates is adjusted.

The rates are adjusted on a per queue level, per subscriber egress agg-rate-limit level, per V-port aggregate level and on per port-scheduler level. Because the ATM termination points are on the BNG, there is a direct view of the encapsulation. The encapsulation information is supplied to the forwarding plane via the control plane.

The rate adjustment are examined in the following commands:

- **frame-based-accounting** under the qos>scheduler-policy hierarchy
- **queue-frame-based-accounting** under the sub-profile>agg-rate-limit hierarchy
- **avg-frame-overhead** under the queue hierarchy
- **encap-offset** under the sub-profile hierarchy

In PPPoA/PPPoEoA scenario there is no last mile rate adjustment (due to the difference in encapsulation) performed. It is assumed that the encapsulation in the last mile and in the intermediate mile is the same. In other words, it is assumed that the DSLAM does not add nor change any encapsulation but instead it only acts as a VPI/VCI cross-connect. If the last mile encapsulation is PPPoA, then the intermediate encapsulation is considered to be PPPoA as well. Similar is valid for PPPoEoA encapsulation.

There are two possible configuration scenarios:

1. The **encap-offset** command in the sub-profile is configured. This command overwrites any other command related to rate conversion that might be configured (**frame-based-accounting**, **queue-frame-based-accounting**, **avg-frame-overhead**, or **avg-frame-size**).
   The **encap-offset** command forces dynamic wire rates calculation in the intermediate mile (directly connected ports) on all levels in the QoS hierarchy. The wire overhead in the intermediate mile takes into account the length of the fixed ATM encapsulation, the variable length of AAL5 encapsulation (including AAL5 48bit boundary padding) and the ATM cellification overhead. The queue stats are also wire based. All calculations are performed in the data plane using the actual packet size. In other words, this command will ensure that the rates on the queue level and the subscriber aggregate level (either through virtual schedulers or egress aggregate-rate limits) are wire based. Port-scheduler and V-port rates are already by default wire based rates and this cannot be changed.
2. The **encap-offset** command is not configured by default. In this case the other rate conversion related commands are in effect (**frame-based-accounting**, **queue-frame-based-accounting**, **avg-frame-overhead**, or **avg-frame-size**). The behavior is the following:
   - Wire rates in the intermediate mile (directly connected ports) are based on the **avg-frame-overhead** command which is provisioned via CLI. If avg-frame-overhead is not provisioned via CLI, by default it is assumed to be 0[%] and the wire rates effectively become data rates (IP payload + IP header + PPP(oE) header + fixed ATM encapsulation).
   - Queue stats (used in accounting) are always ‘data’ stats. This includes the IP Payload + IP header + PPP(oE) header + fixed ATM encapsulation.
   - *agg-rate-limit* (subscriber or Vport) rates are always wire rates (as defined in the first bullet – based on the avg-frame-overhead).
   - Rates in the *port-scheduler-policy* (Vport or physical port) are always wire based (rates (as defined in the first bullet – based on the avg-frame-overhead).
   - The **frame-based-accounting** command under the scheduler-policy will affect rate calculation for virtual schedulers and queues:
     - If this option is configured, the virtual scheduler and queue rates will be wire based.
     - if this option is not configured, the virtual scheduler and queue rates will be data rates
   - *queue-frame-based-accounting* configuration option under the subscriber **agg-rate-limit** command (in sub-profile) affects rate calculations for queues. If this command is configured, the queue rates are wire rates, otherwise they will be data rates.

The **avg-frame-size** command in PPPoA/PPPoEoA is ignored.

Queue rates and subscriber virtual scheduler rates are allowed to be either data rates (one in Figure 104) or on-the-wire-rates (three in Figure 104). Port-scheduler rates, Vport rates and the subscriber agg-rate-limit (in sub-profile) are always on on-the-wire rates.

Figure 104 **ATM Wire Overhead**

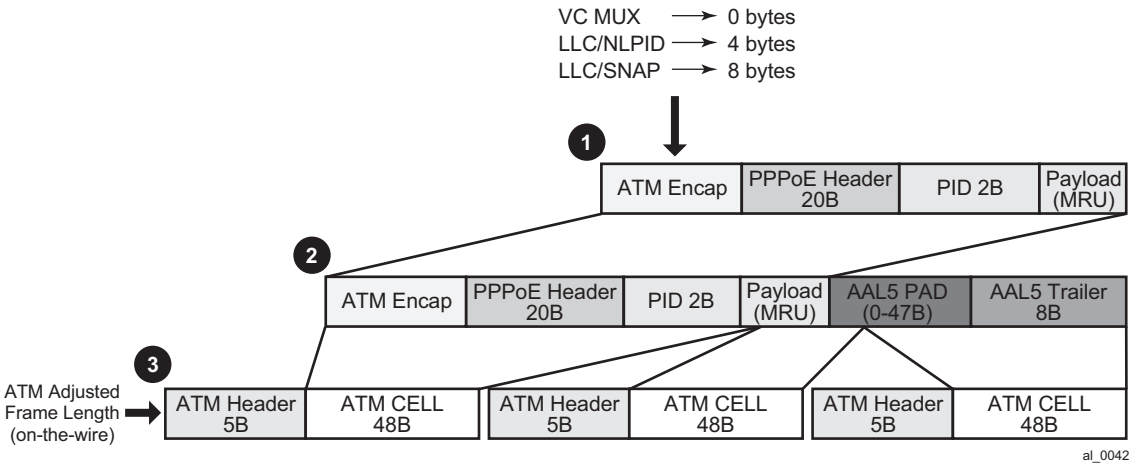

9.3.15.13 Subscriber Instantiation Use Cases

Note that multiple subscribers per single VC are not supported. However, multiple VCs per subscriber are supported.

The following displays examples of how subscriber hosts could be instantiated in PPPoA/PPPoEoA environment.

Case 1

- One host per subscriber
- One ATM VC per subscriber
- Authentication is done via RADIUS or LUDB
- Authentication methods: PAP/CHAP or PADI for PPPoE
- If there is no RADIUS/LUDB authentication, the subscriber-id should be by default the SAP (with VPI:VCI) with default sub-strings. For MSAP this would be configured under the msap-policy, all subs would have the same service in this case (SLA/SUB profiles).

Figure 105 **Subhost per VC**

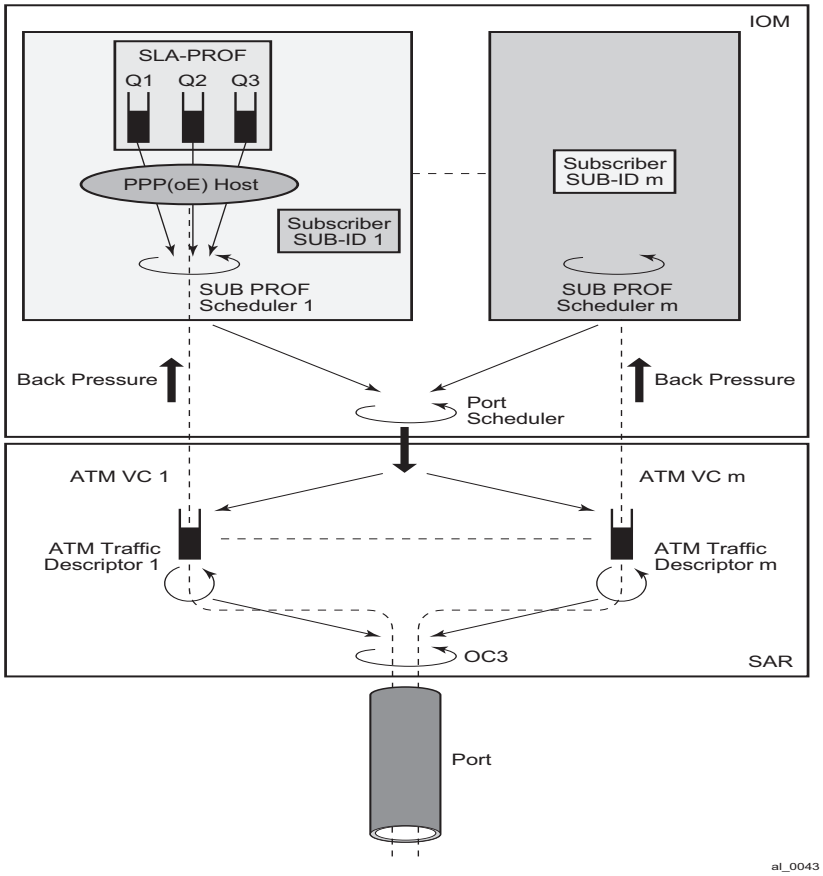

Case 2

- Multiple hosts per subscriber
- Single VC per host
- Multiple VCs per subscriber
- PPPoA or PPPoEoA
- Subhosts for the same subscriber can authenticate using a unique username or the same username

Figure 106 **Multiple VCs per Subscriber**

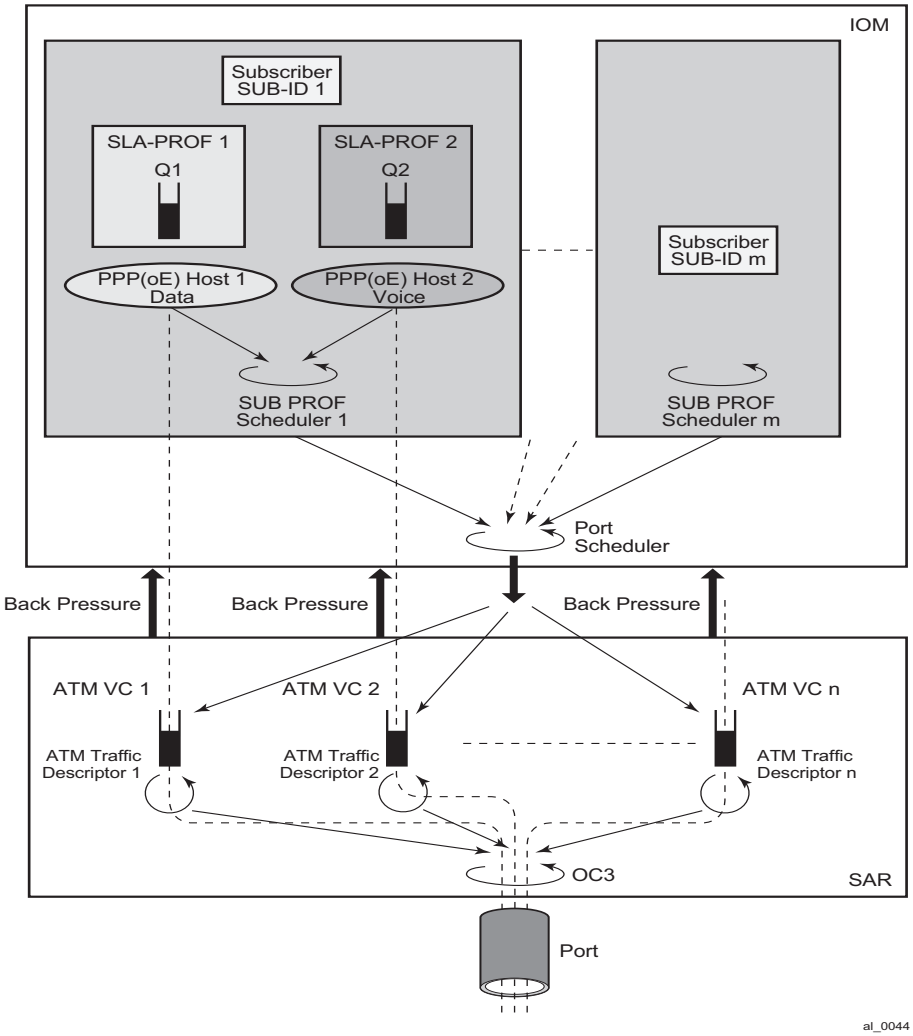

Case 3

- Multiple hosts per subscriber
- Singe VC
- PPPoEoA only
- Per host SLA-PROFILE instantiation

Figure 107 **Multiple Hosts per Subscriber, Single VC**

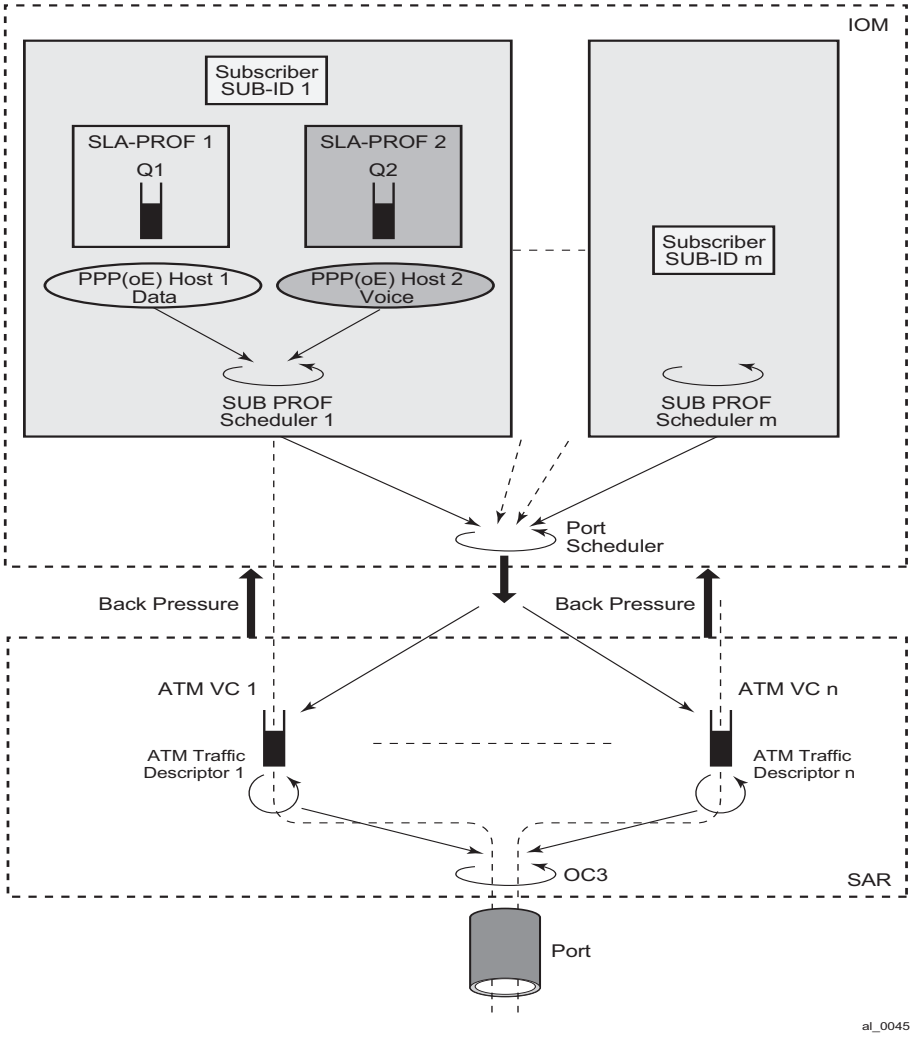

Case 4

- VP Shaping

Figure 108 **VP Shaping**

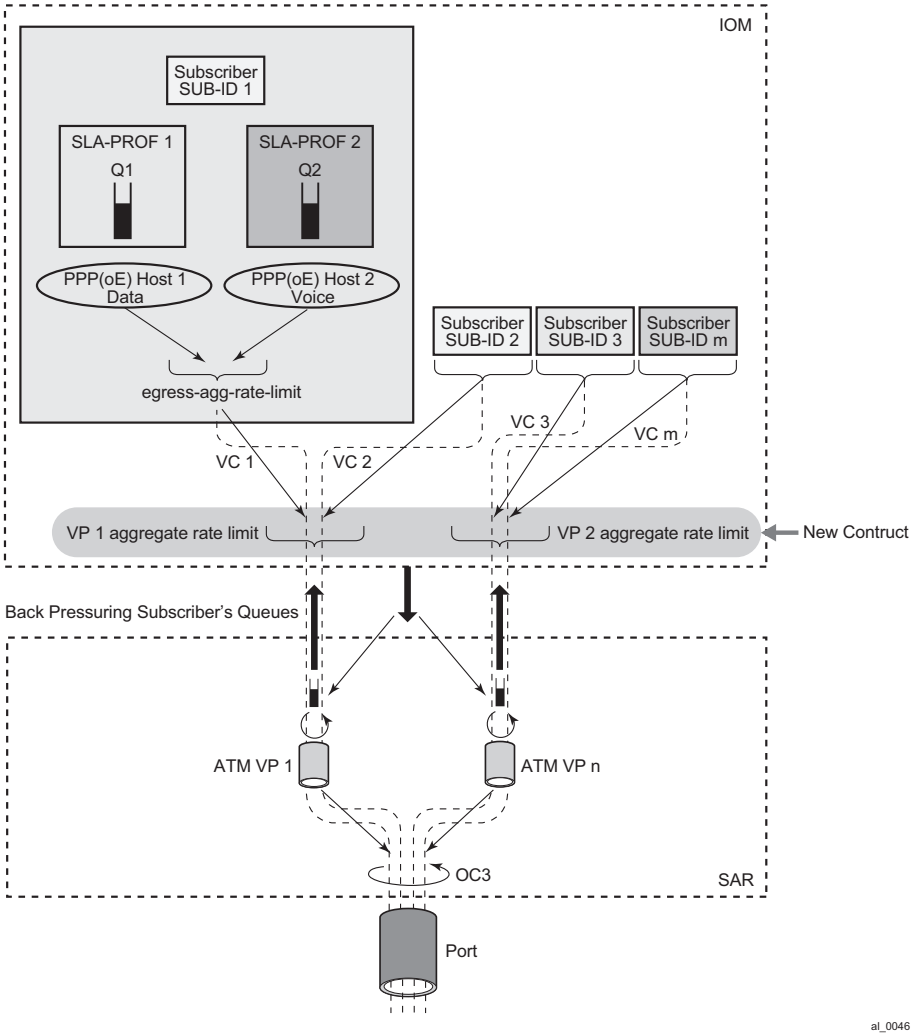

9.3.15.14 Authentication

Authentication for PPPoEoA is the same as in PPPoE:

- based on PADI
- based on PAP/CHAP
- No authentication

Authentication in PPPoA is based on:

- PAP/CHAP

- No Authentication — There is a small percentage of cases where this might be required.

An example of no authentication configuration for PPPoEoA would be:

```
configure
    service ies/vprn
        subscriber-interface <sub-if-name>
            group-interface <grp-if-name>
                dhcp
                    server <ip-address>
                    client-application pppoe  => make PPPoE session the client of
                                                 DHCP server
                    local-dhcp-server <dhcp-server-name>
                    user-db <ludb-name> => allow DHCP server to query LUDB

configure
    subscriber-management
        local-user-db <ludb-name>
            pppoe
                host <host-name>
                host-identification [circuit-id|mac|remote-id|service-name|username]
                options
                address
                identification-strings
```

In the case that there is no PAP/CHAP authentication, a PPPoE host can be identified by a MAC address, circuit/remote-id (inserted by the BNG) or service-name in PPPoE.

9.3.15.15 LUDB Access via Capture SAP

Access to LUDB via a capture SAP is enabled for this feature:

```
configure
    service vpls <id>
        sap <sap-id> capture-sap
            pppoe-user-db <ludb-name>
            ppp-user-db <ludb-name>
```

Note that if the authentication-policy (RADIUS authentication) is specified under the capture SAP, it (the RADIUS authentication) will take precedence over LUDB.

9.3.15.16 Encapsulation Autosensing

As previously discussed, these four static encapsulation types will be supported:

1. aal5nlpid-ppp (LLC/SNAP for PPPoA)
2. aal5snap-bridged (LLC/SNAP for PPPoEoA)
3. aal5mux-ppp (mux PPPoA)
4. aal5-mux-bridged-eth-nofcs (mux PPPoEoA)

These four types of encapsulation are supported for ATM MSAPs as well as for fixed configuration ATM SAPs.

In addition, the LLC/SNAP encapsulation type can be autosensed. This is called autosensing. The keyword for autosensing is **aal5auto**. An option is given to provision encapsulation statically, if needed.

```
sap x/y/z:* capture-sap
atm
encapsulation aal5auto | aal5mux-ppp | aal5nlpid-ppp | aal5mux-pppoe | aal5snap-
bridged-eth-nofcs

sap x/y/z:w/z
atm
encapsulation aal5auto | aal5mux-ppp | aal5nlpid-ppp | aal5mux-pppoe | aal5snap-
bridged-eth-nofcs
```

The aforementioned encapsulation options (including autosensing) is visible only under the SAP hierarchy on group-interfaces and on capture SAPs in VPLS. This is allowed only in 16K-VP mode ATM MDA.

9.3.15.17 SAP Autoprovisioning

In order to simplify the provisioning of the subscriber access ports in the ESM context, a concept similar to managed SAP (MSAP) on Ethernet is introduced. In Ethernet MSAP, an ingress access SAP is automatically created upon receipt of the first (VLAN) tagged packet from the customer side (pending the authentication process).

In our ESM over PPPoA/PPPoEoA case, a number of PVCs is pre-provisioned that will initially be only in a provisioned (or listening, passive) state. This is sometimes referred as a bulk configuration of VC ranges. Once the initial ESM processing in the CPM is completed (for example, the user is authenticated), the ATM SAP is created in the appropriate context (VRF or GRT) and the ATM VC is activated. ATM VC activation means that the ATM VC is associated with a SAP.

The ratio between the maximum number of provisioned VCs vs the maximum number of active VCs on an MDA supporting PPPoA/PPPoEoA is 2:1. This amounts to 32K ATM VCs in the listening state. Out of the total 32K ATM VCs, 16K ATM VCs can be active simultaneously.

There are some differences between the Ethernet MSAP processing and the ATM MSAP processing. On Ethernet there is no need to pre-configure VCs, while on ATM this is necessary due to complexity of ATM layer.

Our current Ethernet based MSAP is configured under VPLS. The same approach is adopted for ATM capture SAP.

Since a combination of ESM over PPPoA/PPPoEoA VCs and other non PPPoA/PPPoEoA VCs on a single physical port is supported, ranges of VCs that will be supported in autoconfiguration are defined. ATM VCs outside this range are available for manual creation. Multiple ranges are necessary in order to address non-contiguous sets of VPI/VCIs.

The following is the CLI syntax:

```
configure
    service vpls <id> customer <customer-id> [create]
        sap x/y/x:*/* capture-sap [create]
            atm
                vc-range <num> vpi-range <vpi-range> vci-range <vci-range>
```

Note that the capture SAP must be configured as '*/*' in place of the VPI/VCI identifiers. Any other combination for the VPI/VCI identifiers is not allowed.

Up to 5 ranges are allowed per capture SAP. VCs configured in this way can carry a single PPP session per VC or multiple PPPoE sessions per VC. Ranges are not allowed unless the ATM MDA is in the 16K-VCs mode.

The total number of VCs that can be fed into a VP is 16K. This includes all VC ranges associated with the VP plus any statically configured VC on the VP.

9.3.15.18 PPP Nodes and ppp-policy

Differences in operation between PPP and PPPoE warrant creation of a new ppp node under the *group-interface* in the CLI that will cover PPP aspects of operation. The existing pppoe node under the *group-interface* is preserved in the CLI. This allows referencing different *user-dbs* for authentication purposes and different session parameters defined in the *ppp-policy* for each session type (pppoe or plain ppp).

PPP node under the *group-interface* is used to cover PPPoA operation while PPPoE node is used to cover PPPoE and PPPoEoA operation. ATM in PPPoEoA is just a transport and as such does not carry any information relevant to PPP operation (like PADx does in PPPoE).

For dynamic SAPs (managed SAPs), the same PPP(oE) related structures are referenced under the capture SAP hierarchy. Under the capture SAP there are no ppp/pppoe nodes (like they are under the group-interface hierarchy). In order to differentiate between *ppp* and *pppoe* clients, the **ppp-policy** command is introduced in addition to the **pppoe-policy** command. The **ppp-policy** under the capture SAP is needed for the definition of session parameters before the **group-interface** (where normally session parameters are referenced) is determined.

Two commands for LUDB access are available under the same **capture-sap** hierarchy (**pppoe-user-db** and **ppp-user-db**).

The **ppp-policy** under the **subscr-mgmt** hierarchy contains PPP and PPPoE session parameters. PPP parameters are applicable to both session types (PPP and PPPoE) while PPPoE parameters are applicable only to PPPoE session type. The PPPoE parameters are ignored for PPP sessions.

9.3.15.19 MTU Considerations

MRU configuration option negotiated during the LCP phase in PPPoA is based on the following command:

```
config>subscr-mgmt>ppp-policy#
 ppp-mtu
```

By default, this command is disabled and consequently will negotiate the default MRU of 1500B, as long as the ATM port's MTU can accommodate at least 1500B:

```
config>port>sonnet-sdh>path#
mtu
```

The MRU option in PPPoA refers to the PPP packet length (PID+Information+Padding) that is ATM encapsulated.

9.3.15.19.1 PPP(oE) Session Antispoofing

Antispoofing filters need to be in place in order to prevent the hijacking of a PPPoA/ PPPoEoA session. For successful anti-spoofing, the following fields are accessible:

- Source IP address (PPPoA and PPPoEoA)
- VPI/VCI pair which is equivalent to the SAP (PPPoA/PPPoEoA)
- Source MAC Address (PPPoEoA only)
- session-ID (PPPoEoA only)

For locally terminated PPPoEoA subscriber hosts, access to all fields are required (a, b, c and d). Antispoofing is defined under the following hierarchy:

```
config
    service <svc-name>
        subscriber-interface <sub-if-name>
            group-interface <grp-if-name>
                pppoe
                    anti-spoof [mac-sid | mac-sid-ip]
```

This behavior matches our current PPPoE behavior. The default antispoofing option is set to mac-sid, which means that the incoming traffic is checked against the source MAC address and the session-ID. Antispoofing under this hierarchy cannot be disabled.

The source IP address can be added to the source MAC and session-ID (mac-sid-ip).

Note that the group-interface is the lowest granularity at which options can be enabled. In other words, these two options (mac+sid or mac+sid+ip) cannot be changed at the SAP level.

The VPI/VCI pair is always be checked against the incoming traffic, regardless of the configuration option as the VPI/VCI pair is an intrinsic part of the SAP (similar to VLAN tags on Ethernet).

In case that a subscriber host is a routed host (managed routes), the nh-mac antispoofing option must be enabled under the managed sap (msap-policy) or **group-interface->sap** level. Otherwise, managed routes for the host would **not** be installed. The nh-mac option forces a lookup of the incoming packet based on the mac+sid of the originally created host (CPE device). Note that only the **group-interface->sap>anti-spoof** and msap-policy hierarchies contain the **nh-mac** option, and not the pppoe->anti-spoof hierarchy.

For locally terminated PPPoA subscriber hosts, only access to source IP address is available (src mac and session-ID fields are non-existent in PPP).

The default antispoofing is based on the VPI/VCI pair + IP. This cannot be changed by configuration, except when the session has managed routes or is a LAC session. In the case of managed routes (routed host) a lookup will be done based on the SAP (VPI/VCI) only. The anti-spoofing for PPPoA should be set to anti-spoof nh-mac, under msap-policy or grp-if>sap hierarchy even though PPPoA has no MAC address.

For LAC, the behavior is the following:

- For PPPoE traffic, antispoofing is always done based only on the SAP+mac+session-id.
- For PPP traffic, antispoofing is always based on VPI/VCI pair (SAP only). The IOM automatically determines whether IP based antispoofing can be done (e.g. no IP based antispoofing for LAC and managed routes).

There are two other nodes on which antispoofing can be configured:

```
configure
    service ies/vprn
        subscriber-interface <sub-if-name>
            group-interface <grp-if-name>
                sap <sap-id>
                    anti-spoofing [ip|ip-mac|nh-mac]
```

A few points in regards to the above hierarchy:

- **ip** - This option can be only used in BSM mode. This is used only for IPoE.
- **ip-mac** – Lookup is performed based on the combination of the IP address and MAC address (incipient host). This is used for IPoE. For PPPoE, it will be overwritten by the configuration option under the group-interface-> ppp node (mac+sid or mac-sid+ip).
- **nh-mac** – Lookup is based on the MAC only for IPoE or mac+sid for PPPoE.

```
configure
    subscriber-management
        msap-policy <msap-pol-name>
            ies-vprn-only-sap-parameters
                anti-spoof [ip-mac|nh-mac]
```

This is similar to the previous case (under the group-interface>sap hierarchy) with the exception that the pure IP option is not supported under the MSAP. The reason is that the IP option can only be used in BSM (under the SAP node) whereas MSAP can only be used in ESM (PPPoX).

Note that for IPoE, the entire control over antispoofing under the (M)SAP node, while for PPPoE, the anti-spoofing control is distributed between the (M)SAP node and the **group-interface>ppp** node.

9.3.16 Multi-Chassis Synchronization

Figure 109 shows the configuration under which synchronization of subscriber management information is performed. As depicted, a single access node aggregating several subscriber lines is dual- homed to redundant-pair of nodes.

Figure 109 **Dual-Homing Configuration**

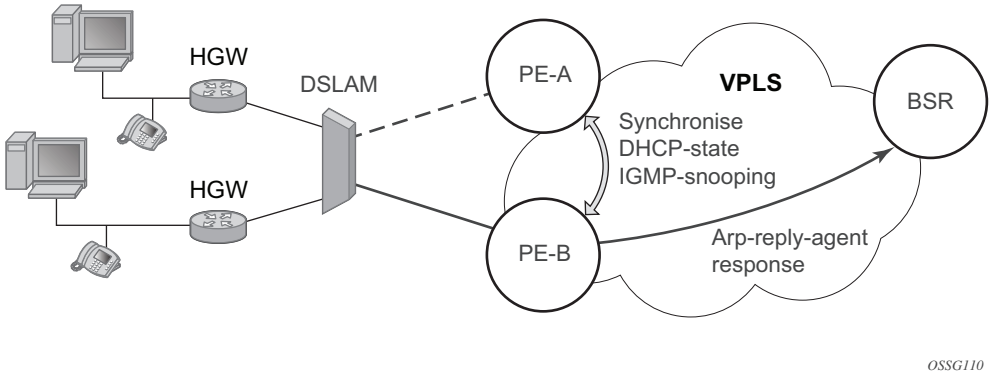

Enabling subscriber management features (whether basic subscriber-management (BSM) or enhanced subscriber management (ESM)) causes the node to create and maintain state information related to a given subscriber-host. This information is synchronized between redundant-pair nodes to secure non-stop service delivery in case of the switchover.

9.3.16.1 Overview

The synchronization process provides the means to manage distributed database (the Multi-Chassis Synchronization (MCS) database), which contains the dynamic state information created on any of the nodes by any application using its services. The individual entries in the MCS database are always paired by peering-relation, sync-tag and application-id. At any time the given entry is related to the single redundant-pair objects (two SAPs on two different nodes) and hence stored in a local MCS database of the respective nodes.

Internally, peering-relation and sync-tag are translated into a port and encapsulation value identifying the object (SAP) that the given entry is associated with. The application-id then identifies the application which created the entry on one of the nodes. There are three basic operations that the application can perform on MCS database. The MCS database will always synchronize these operations with its respective peer for the given entry.

The following principles apply:

- add-operation — Any dynamic-state created in the application is pushed to the MCS database. MCS then creates and synchronizes with the corresponding peer provided (if configured). The application in the peer node is then notified as soon as the entry has been created. Similarly, the application in the local node (the node where the state has been created) is notified that entry has been synchronized (MCS is “in-sync” state). This operation will be also used to modify existing MCS database entry.
- local-delete — The MCS database entry is marked as no longer in use locally and this information is sent to the peer node. If the information is no longer used by applications on both nodes (the application in remote-node has already issued local-delete before), it is removed from database.
- global-delete — The MCS database entry is removed from both nodes and from the application in the remote node.

The choice of the operation in corresponding situation is driven by the application. The following general guidelines are observed:

- An event which leads to a dynamic-state deletion on a standby chassis will be handled as “local-delete”.
- An event which leads to a dynamic-state deletion on an active chassis will be handled as “global-delete”.
- An exception to above the rules is an explicit **clear** command which will be handled as global-delete regardless of where the command was executed.

As previously stated, the MCS process automatically synchronizes any database operation with the corresponding peer. During this time, the MCS process maintains state per peer indicating to the applications (and network operator) the current status, such as in-sync, synchronizing or sync_down. These states are indicated by corresponding traps.

9.3.16.1.1 Loss of Synchronization and Reconciliation

Each time the connection between the redundant pair nodes is (re)established the MCS database will be re-synchronized. There are several levels of connectivity loss which may have different effect on amount of data being lost. To prevent massive retransmissions when the synchronization connection experiences loss or excessive delay, the MCS process implementation will take provisions to ensure following:

- In the case of a reboot of one or both nodes or establishing the peering for the first time, the full MCS database will be reconciled.

- In the case that the MCS communication is lost and then re-established but neither node rebooted during the connection loss, only the information not synchronized during this time will be reconciled (using sequence numbers helps identify information which was not synchronized).
- In the case that MCS communication is lost because of excessive delay in ACK messages but no information has been effectively lost, the MCS process indicates a loss of synchronization but no reconciliation is performed.

9.3.17 Subscriber Routed Redundancy Protocol (SRRP)

9.3.17.1 SRRP Messaging

SRRP uses the same messaging format as VRRP with slight modifications. The source IP address is derived from the system IP address assigned to the local router. The destination IP address and IP protocol are the same as VRRP (224.0.0.18 and 112, respectively).

The message type field is set to 1 (advertisement) and the protocol version is set to 8 to differentiate SRRP message processing from VRRP message processing.

The vr-id field has been expanded to support an SRRP instance ID of 32 bits.

Due to the large number of subnets backed up by SRRP, only one message every minute carries the gateway IP addresses associated with the SRRP instance. These gateway addresses are stored by the local SRRP instance and are compared with the gateway addresses associated with the local subscriber IP interface.

Unlike VRRP, only two nodes may participate in an SRRP instance due the explicit association between the SRRP instance group IP interface, the associated redundant IP interface and the multi-chassis synchronization (MCS) peering. Since only two nodes are participating, the VRRP skew timer is not utilized when waiting to enter the master state. Also, SRRP always preempts when the local priority is better than the current master and the backup SRRP instance always inherits the master's advertisement interval from the SRRP advertisement messaging.

SRRP advertisement messages carry a *becoming-master* indicator flag. The *becoming-master* flag is set by a node that is attempting to usurp the master state from an existing SRRP master router. When receiving an SRRP advertisement message with a better priority and with the *becoming-master* flag set, the local master initiates its *becoming-backup* state, stops routing with the SRRP gateway MAC and sends an SRRP advertisement message with a priority set to zero. The new master continues to send SRRP advertisement messages with the *becoming-*

master flag set until it either receives a return priority zero SRRP advertisement message from the previous master or its *becoming-master* state timer expires. The new backup node continues to send zero priority SRRP advertisement messages every time it receives an SRRP advertisement message with the *becoming-master* flag set. After the new master either receives the old master's priority zero SRRP advertisement message or the *become-master* state timer expires, it enters the *master* state. The *become-master* state timer is set to 10 seconds upon entering the *become-master* state.

The SRRP advertisement message is always evaluated to see if it has higher priority than the SRRP advertisement that would be sent by the local node. If the advertised priority is equal to the current local priority, the source IP address of the received SRRP advertisement is used as a tie breaker. The node with the lowest IP address is considered to have the highest priority.

The SRRP instance maintains the source IP address of the current master. If an advertisement is received with the current masters source IP address and the local priority is higher priority than the masters advertised priority, the local node immediately enters the *becoming-master* state unless the advertised priority is zero. If the advertised priority is zero, the local node bypasses the *becoming-master* state and immediately enters the *master* state. Priority zero is a special case and is sent when an SRRP instance is relinquishing the master state.

9.3.17.2 SRRP and Multi-Chassis Synchronization

In order to take full advantage of SRRP resiliency and diagnostic capabilities, the SRRP instance should be tied to a MCS peering that terminates on the redundant node. The SRRP instance is tied to the peering using the **srrp** *srrp-id* command within the appropriate MCS peering configuration. Once the peering is associated with the SRRP instance, MCS will synchronize the local information about the SRRP instance with the neighbor router. MCS automatically derives the MCS key for the SRRP instance based on the SRRP instance ID. For example, an SRRP instance ID of 1 would appear in the MCS peering database with a MCS-key srrp-0000000001.

The SRRP instance information stored and sent to the neighbor router consists of:

- The SRRP instance MCS key
- Containing service type and ID
- Containing subscriber IP interface name
- Subscriber subnet information
- Containing group IP interface information

- The SRRP group IP interface redundant IP interface name, IP address and mask
- The SRRP advertisement message SAP
- The local system IP address (SRRP advertisement message source IP address)
- The Group IP interface MAC address
- The SRRP gateway MAC address
- The SRRP instance administration state (up or down)
- The SRRP instance operational state (disabled/becoming-backup, becoming-master, master)
- The current SRRP priority
- Remote redundant IP interface availability (available or unavailable)
- Local receive SRRP advertisement SAP availability (available or unavailable)

9.3.17.3 SRRP Instance

The SRRP instance uses the received information to verify provisioning and obtain operational status of the SRRP instance on the neighboring router.

9.3.17.3.1 SRRP Instance MCS Key

The SRRP instance MCS key ties the received MCS information to the local SRRP instance with the same MCS key. If the received key does not match an existing SRRP instance, the MCS information associated with the key is ignored. Once an SRRP instance is created and mapped to an MCS peering, the SRRP instance evaluates received information with the same MCS key to verify it corresponds to the same peering. If the received MCS key is on a different peering than the local MCS key an SRRP peering mismatch event is generated detailing the SRRP instance ID, the IP address of the peering the MCS key is received on and the IP address to which the local MCS key is mapped. If the peering association mismatch is corrected, an SRRP peering mismatch clear event is generated.

9.3.17.3.2 Containing Service Type and ID

The Containing Service Type is the service type (IES or VPRN) that contains the local SRRP instance. The Containing Service ID is the service ID of that service. This information is supplied for troubleshooting purposes only and is not required to be the same on both nodes.

9.3.17.3.3 Containing Subscriber IP Interface Name

The containing subscriber IP interface name is the subscriber IP interface name that contains the SRRP instance and its group IP interface. This information is supplied for troubleshooting purposes only and is not required to be the same on both nodes.

9.3.17.3.4 Subscriber Subnet Information

The subscriber subnet information includes all subscriber subnets backed up by the SRRP instance. The information for each subnet includes the Owned IP address, the mask and the gateway IP address. If the received subscriber subnet information does not match the local subscriber subnet information, an SRRP Subscriber Subnet Mismatch event is generated describing the SRRP instance ID and the local and remote node IP addresses. Once the subscriber subnet information matches, an SRRP Subscriber Subnet Mismatch Clear event is generated.

9.3.17.3.5 Containing Group IP Interface Information

The containing group IP interface information is the information about the group IP interface that contains the SRRP instance. The information includes the name of the group IP interface, the list of all SAPs created on the group IP interface, the administrative and operational state of each SAP and the MCS key and the peering destination IP address associated with each SAP. To obtain the MCS information, the SRRP instance queries MCS to determine the peering association of the SRRP instance and then queries MCS for each SAP on the group IP interface. If the local SRRP instance is associated with a different MCS peering than any of the SAPs or if one or more SAPs are not tied to an MCS peering, an SRRP group interface SAP peering mismatch event is generated detailing the SRRP instance ID, and the group IP interface name.

When receiving the remote containing group IP interface information, the local node compares the received SAP information with the local group IP interface SAP information. If a local SAP is not included in the SAP information or a remote SAP is not included in the local group IP interface, an SRRP Remote SAP mismatch event is generated detailing the SRRP instance ID and the local and remote group IP interface names. If a received SAP’s MCS key does not match a local SAP's MCS Key, an SRRP SAP MCS key mismatch event is generated detailing the SRRP instance ID, the local and remote group IP interface names, the SAP-ID and the local and remote MCS keys.

9.3.17.3.6 Remote Redundant IP Interface Mismatch

If the group IP remote redundant IP interface address space does not exist, is not within the local routing context for the SRRP instances group IP interface or is not on a redundant IP interface, the local node sends redundant IP interface unavailable to prevent the remote neighbor from using its redundant IP interface. An SRRP redundant IP interface mismatch event is generated for the SRRP instance detailing the SRRP instance, the local and remote system IP addresses, the local and remote group IP interface names and the local and remote redundant IP interface names and IP addresses and masks. The local redundant IP interface may still be used if the remote node is not sending redundant IP interface unavailable.

9.3.17.3.7 Remote Sending Redundant IP Interface Unavailable

If the remote node is sending redundant IP interface unavailable, the local node will treat the local redundant IP interface associated with the SRRP instances group IP interface as down. A Local Redundant IP Interface Unavailable event is generated detailing the SRRP instance ID, the local and remote system IP addresses, the local group IP interface name, the local redundant IP interface name and the redundant IP interface IP address and mask.

9.3.17.3.8 Remote SRRP Advertisement SAP Non-existent

If the remote node’s SRRP advertisement SAP does not exist on the local SRRP instances group IP interface, the local node sends local receive SRRP advertisement SAP unavailable to the remote node. An SRRP receive advertisement SAP non-existent event is generated detailing the SRRP instance ID, the local and remote system IP addresses, the local group IP interface name and the received remote SRRP advertisement SAP. Since SRRP advertisement messages cannot be received, the local node will immediately become master if it has the lower system IP address.

9.3.17.3.9 Remote Sending Local Receive SRRP Advertisement SAP Unavailable

If the local node is receiving local receive SRRP advertisements stating that the SAP is unavailable from the remote node, an SRRP Remote Receive advertisement SAP Unavailable event will be generated. This details the SRRP instance ID, the local and remote system IP addresses, the remote group IP interface name and the local SRRP advertisement SAP. Since the remote node cannot receive SRRP advertisement messages, the local node will immediately become master if it has the lower system IP address.

9.3.17.3.10 Local and Remote Dual Master Detected

If the local SRRP state is master and the remote SRRP state is master, an SRRP dual master event is generated detailing the SRRP instance ID and the local, remote system IP addresses and the local and remote group IP interface names and port numbers.

9.3.17.4 Subscriber Subnet Owned IP Address Connectivity

In order for the network to reliably reach the owned IP addresses on a subscriber subnet, the owning node must advertise the IP addresses as /32 host routes into the core. This is important since the subscriber subnet is advertised into the core by multiple routers and the network will follow the shortest path to the closest available router which may not own the IP address if the /32 is not advertised within the IGP.

9.3.17.5 Subscriber Subnet SRRP Gateway IP Address Connectivity

The SRRP gateway IP addresses on the subscriber subnets cannot be advertised as /32 host routes since they may be active (master) on multiple group IP interfaces on multiple SRRP routers. Without a /32 host route path, the network will forward any packet destined to an SRRP gateway IP address to the closest router advertising the subscriber subnet. While a case may be made that only a node that is currently forwarding for the gateway IP address in a master state should respond to ping or other diagnostic messages, the distribution of the subnet and the case of multiple masters make any resulting response or non-response inconclusive at best. To provide some ability to ping the SRRP gateway address from the network side reliably, any node receiving the ICMP ping request will respond if the gateway IP address is defined on its subscriber subnet.

9.3.17.6 Receive SRRP Advertisement SAP and Anti-Spoof

The group IP interface SAPs are designed to support subscriber hosts and perform an ingress anti-spoof function that ensures that any IP packet received on the group IP interface is coming in the correct SAP with the correct MAC address. If the IP and MAC are not registered as valid subscriber hosts on the SAP, the packet is silently discarded. Since the SRRP advertisement source IP addresses are not subscriber hosts, an anti-spoof entry will not exist and SRRP advertisement messages would normally be silently discarded. To avoid this issue, when a group IP interface SAP is configured to send and receive SRRP advertisement messages, anti-spoof processing on the SAP is disabled. This precludes subscriber host management on the SRRP messaging SAP.

9.3.18 PPPoE MC Redundancy

This feature minimizes the downtime for PPPoE clients in an ESM environment when a single node fails.

But it is not necessary that an entire BNG fails before it triggers the corrective action. The solution outlined in this document will natively include protection against interfaces and line card failures within the BNG. The redundant (protective) entity, however, does not reside within the same BNG on which the failure occurs but instead it is on a separate BNG node.

The PPPoE MC Redundancy is based on SRRP and MC-LAG because SRRP is already established in ESM providing IPoE MC Redundancy. With some modifications, SRRP approach is adopted to PPPoE deployments.

9.3.18.1 SRRP Considerations for PPPoE

SRRP is based on VRRP whose purpose is to provide a default gateway redundancy for clients sharing the transport medium such as Ethernet. IPoE would be a typical example of this where IPoE clients use a virtual IP and MAC address that is shared between two default gateway nodes in the Master/Backup configuration. SRRP supports only two nodes in a cluster but VRRP allows multiple nodes to be configured in a cluster with a priority that will determine which node will assume Mastership. Although it is mandatory for the proper operation of IPoE clients that the same SRRP IP address is shared between the two BNG nodes providing

redundancy, having the same SRRP IP address is not necessary for the operation of SRRP itself. In other words, SRRP itself (Master or Backup states) works with different SRRP IP addresses on each node. Same is valid for MAC addressing. It is possible by configuration that the redundant BNG nodes use different IP/MAC addresses on a pair of SRRP instances.

Upon a switchover, a gratuitous ARP is sent from a newly selected active node so that each IPoE client can update the ARP table, if the MAC address has indeed changed (it does not have to). More importantly, if an Layer 2 aggregation network is in place between the BNG and the IPoE client, all intermediate Layer 2 devices will have to update their port-to-mac mappings (Layer 2 FDB). The above described process will ensure proper packet addressing on the IPoE client side as well as the proper forwarding path through Layer 2 aggregation network to the newly activated BNG.

When considering PPPoE in conjunction with SRRP, keep in mind that PPP protocol (point-to-point protocol) is adopted for the Ethernet (shared medium) by enabling an extra Ethernet related layer in PPP that allows sharing of point-to-point sessions over Ethernet (shared medium). The result is a PPPoE protocol designed to ‘tunnel’ each PPP session over Ethernet.

PPPoE is not aware of ARP (Address Resolution Protocol) and it will not react to gratuitous ARP packets sent by a newly active BNG. The destination MAC address that PPPoE clients will use when sending traffic is determined not by ARP but by the PPPoE Discovery phase at the beginning of the session establishment. This originally discovered destination MAC is used throughout the lifetime of the session. This has a couple of consequences:

1. If SRRP is used for PPPoE then the ‘SRRP’ MAC address between the redundant BNG nodes must be shared. It is not allowed to use a unique ‘SRRP’ MAC address per BNG in the redundant pair of BNG nodes (as it is allowed today for IPoE). Every PADx conversation is based on the SRRP shared MAC address, that is, the PADO reply must have the shared SRRP MAC address as the source MAC. This has a significant impact on the operation of MSAP in conjunction with this feature.
2. Since PPPoE sessions are not ARP aware, the only purpose of the gratuitous ARP would be to update the Layer 2 FDB in the aggregation network (and not the PPPoE client destination MAC address). For IPoE, the gratuitous ARP is sent for all subnet gateway IP addresses found under the subscriber interface over either all SAPs (default) or top-tags only. For PPPoE, the gratuitous ARP is sent only for the system IP address. The purpose of the gratuitous ARP in PPPoE scenario is only to update Layer 2 network path which is otherwise IP

unaware. It is not necessary to send the gratuitous ARP for every default-gateway address found under the subscriber-interface. Since this feature is only applicable to PPPoE deployments, therefore, only PPPoE is present under the group interface. This is indicated by the following command under the SRRP node:

```
group-interface <name>
    srrp <id>
        one-garp-per-sap
```

9.3.18.1.1 SRRP Fact-Checks

1. Once Multi-chassis Synchronization (MCS) for subscriber management and SRRP is enabled, both BNG nodes, Master and Backup will in general forward packets (for subscribers) in both directions.
2. Traffic flows through an SRRP enabled node according to the entries in the SRRP sync database and the SRRP state of the node:
   - Backup SRRP directs downstream traffic over the redundant-interface towards the Master SRRP node. If the redundant interface is unavailable, traffic is sent directly to the subscriber.
   - Master SRRP always directly forwards the downstream traffic towards the subscriber.
   - In the upstream direction, the active SRRP node accepts subscriber traffic addressed either to the MAC address of the SRRP active group OR the native interface MAC address.
   - The standby node accepts in the upstream direction only packets addressed to its native interface MAC address.
3. If both SRRP nodes become Masters then both forward traffic to/from subscribers unaware of the link failure somewhere in the Layer 2 network. As a result, downstream traffic can be blackholed. Whether downstream traffic will be lost depends on the native routing on the network side, which is unaware of the failures in the aggregation network.

9.3.18.2 State Synchronization

PPPoE sessions are synchronized between the redundant BNG nodes. The subscriber synchronization is achieved through Multi-Chassis Synchronization (MCS) protocol in a similar way it is performed for IPoE.

```
multi-chassis
           peer <IP@>create
               sync
                   local-dhcp-server
```

```
                    SRRP
                    sub-mgmt [ipoe | pppoe]
                :
:
                    no shutdown
                exit
                no shutdown
            exit
```

Two keywords, **ipoe** and **pppoe** enable a more granular control over which type of subscribers the MCS should be enabled.

Subscriber synchronization is important for following reasons:

1. Forwarding of downstream traffic between the redundant BNG nodes through a redundant interface is an artifact of how natural routing steers traffic through the network.
2. Subscriber instantiation on the node which did not originally create subscriber session. This drastically reduces downtime during the SRRP switchover.
3. Monitors operational aspects of the subscriber management through show commands.

9.3.18.2.1 PPPoE Multi-chassis Synchronization (MCS) Model

PPPoE MCS model is based on SRRP synchronization and can be used in a centralized or distributed environment with or without Layer 2 aggregation network in-between DSLAMs and BNG nodes. The failure detection speed is dependent on SRRP timers. Traffic load can be balanced per SRRP group over the two links. In this model (Figure 110), PPPoE states are synchronized between the redundant BNG nodes. If one BNG fails, the newly activated BNG sends out a ‘MAC update’ (gratuitous ARP) message prompting the intermediate Layer 2 nodes to update their forwarding tables so that forwarding can resume. The SRRP timers can be configured in the sub-second range. In reality, the limiting factor for timer values is the scale of the deployment, in particular the number of SRRP groups per node.

Figure 110 **Fully Redundant "Statefull 1:1" Model**

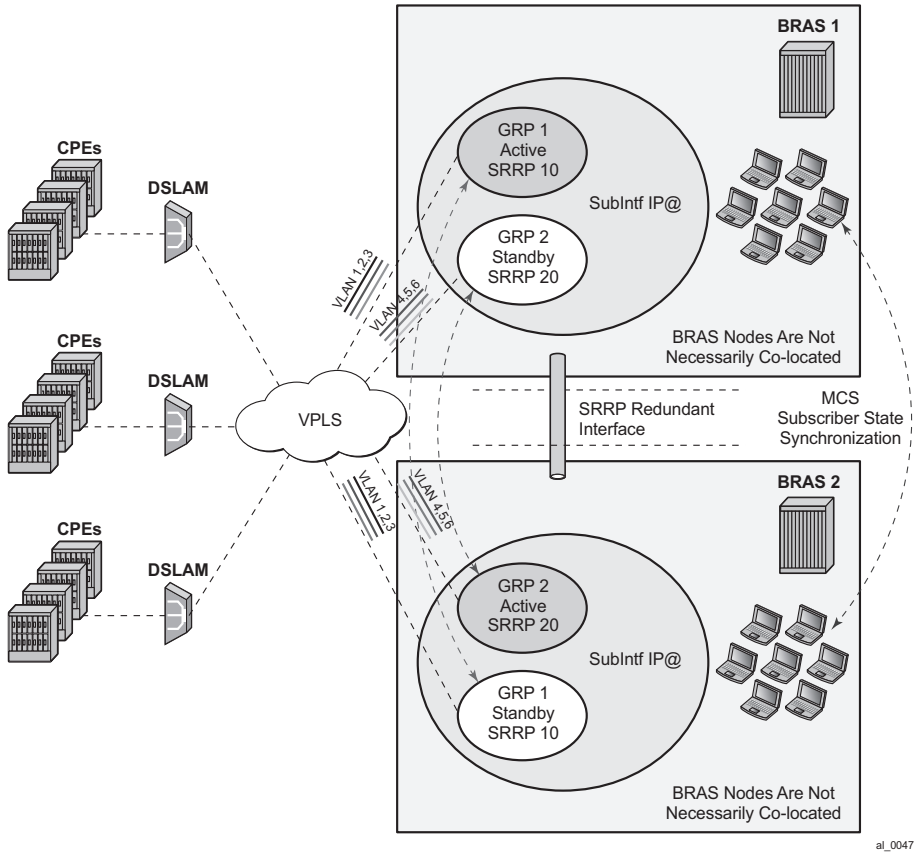

9.3.18.3 Traffic Control and Redundant Interface

To preserve QoS and Accounting, subscriber's traffic must flow in both directions through the Master BNG node.

In the upstream direction, this is always true as traffic is steered to the master SRRP node just by the virtue of SRRP operation.

In the downstream direction which represents bulk of traffic, SRRP can not be relied up on to steer traffic through the Master node. This poses a problem in a very common environment where IP subnets are shared over multiple group-interfaces with SRRP enabled. A particular subnet will be advertised to the network side from both BNG nodes, Master and Backup. Natural routing on the network side determines which BNG node receives subscriber's traffic in the downstream direction. If the Backup SRRP node receives the traffic, it cannot simply send the traffic directly to the access network where the subscriber resides by just inserting

the source MAC address of the SRRP instance in the outgoing packet. This would break the operation of SRRP. Instead, the Backup BNG node must send the traffic to the Master BNG node via a redundant-interface. The Master SRRP node would then forward traffic directly to the subscriber. Source MAC address of this traffic would then be the MAC address of SRRP instance. This traffic shunting over the redundant interface can result in a substantial load on the link between the two BNG nodes.

The increase in shunted traffic can quickly become an issue if the redundant BNG nodes that are not collocated. To minimize the shunt traffic, more granular routing information must be presented to the network core. This would lead to more optimal routing where downstream subscriber traffic would be directed towards the Master BNG node, without the need to cross the redundant interface. The downside of this approach is that this would further fragment the IP address space within the network core. In the extreme case where /32 (subscriber) IP addresses are advertised, the churn that /32s can cause in the core routing would most likely be unsustainable. In this case, routing updates in the core would be triggered by subscribers coming on/off-line.

Optimal operation would call for the shunt traffic to be eliminated and at the same time, a high IP route aggregation on the network side is achieved. The existence of the shunt traffic stems from the fact that routing protocols advertise subscriber subnets into the network with no awareness of the SRRP activity state (Master/Standby). To address this problem along with better aggregation of advertised subnets, two SRRP enhancements are introduced:

- SRRP fate-sharing
- SRRP aware routing

Both of these concepts are described in SRRP Enhancement.

Traffic destined to or from the subscriber is forwarded under the condition that the subscriber-interface is operationally UP. This applies also to shunting of downstream subscriber traffic from the STANDBY to MASTER node. It is always necessary to keep the subscriber-interface operationally UP by configuring a dummy group-interface with a **oper-up-while-empty** command under it. This is especially true for the MC-LAG which causes the messaging SAP on the STANDBY node always to be in the INIT state. In case that MSAPs are used on such group-interfaces, the group-interfaces would be also operationally DOWN, causing the subscriber-interface to be operationally DOWN.

9.3.18.3.1 Subnet Assignment and Advertisement - Option 'A'

A single IP subnet is used for all subscribers terminated within the redundant BNG nodes. The upside of the Option 'A' is that it offers aggregated IP addressing in the network core per pair of redundant BNG nodes. The downside is that the subscriber termination point (active BNG for the SRRP group) is hidden from the network core. Since both BNG nodes share the same IP subnet for the subscribers, the natural routing can cause downstream traffic to be sent to the standby BNG which in turn will have to shunt the traffic to the active BNG. It is likely that half of the traffic will be shunted over the redundant-interface with this approach. This scenario is shown in Figure 111.

Figure 111 **Shared Subscriber IP Space**

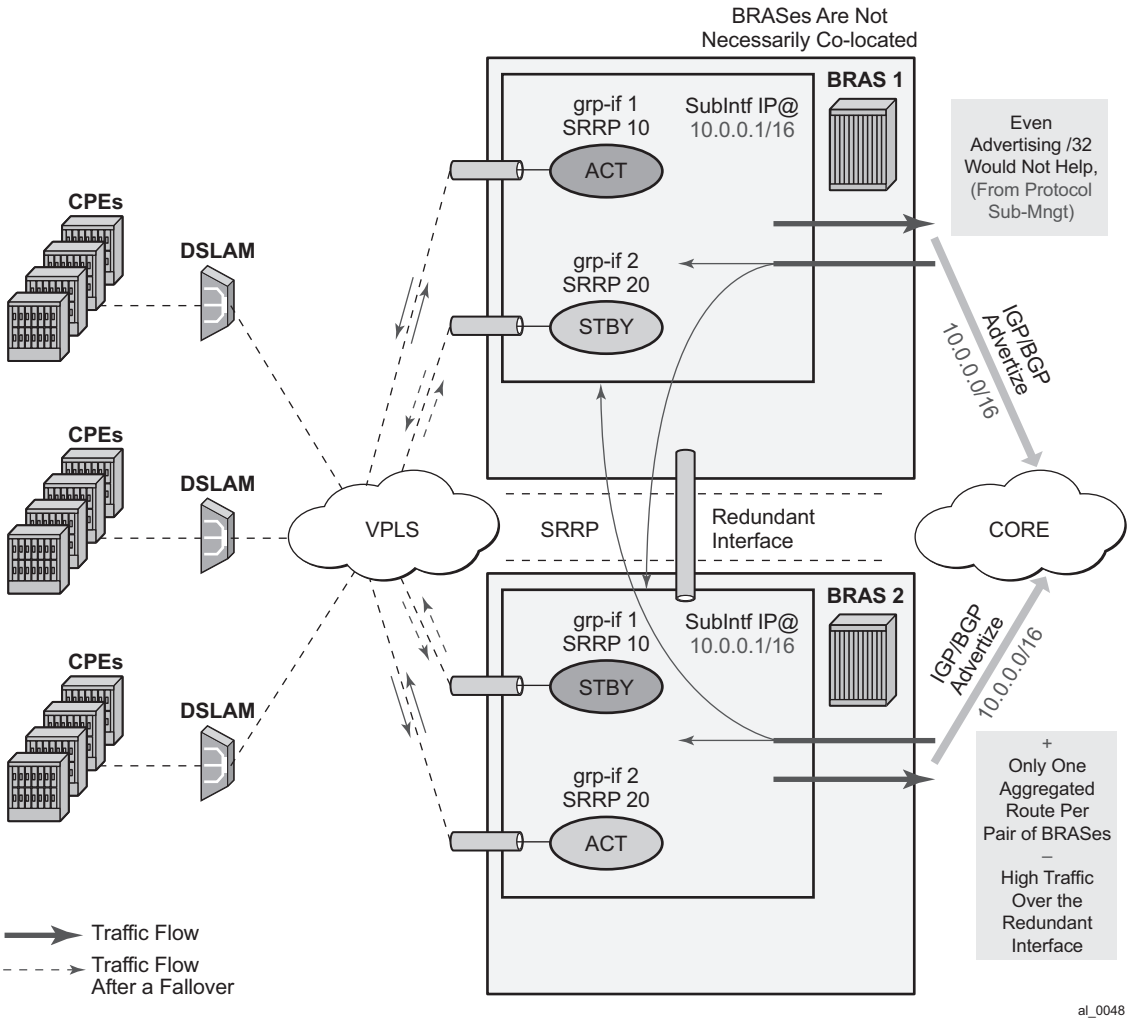

9.3.18.3.2 Subnet Assignment and Advertisement - Option ‘B’

With the option ‘B’, an IP address pool (or subnet) can be allocated per group of SRRP instances that are in the Master state. The routing decision on the network side is further influenced by the static increase of the metric of the advertised route on the BNG node hosting the active SRRP groups (Figure 112).

This approach would cause greater IP space segmentation in the network core, but at the same time, it would indirectly provide more information about the subscriber whereabouts and thus minimize or eliminate the shunt traffic during the normal operation. However, in the case of a SRRP switchover, the shunt traffic would ensue. The amount of the shunted traffic would depend on the scale of the failure. From the Figure 112, it can be concluded that:

- In the depicted scenario on Figure 112, there is no shunted traffic.
- If any of the SRRP instances transitions out of the Master state, traffic for an entire IP network associated with this failed SRRP instance would be shunted. The reason for this is that the advertised route metric is static and it does not follow changes in SRRP state.

Figure 112 **Option 'B' – IP Subnet per Active SRRP Group**

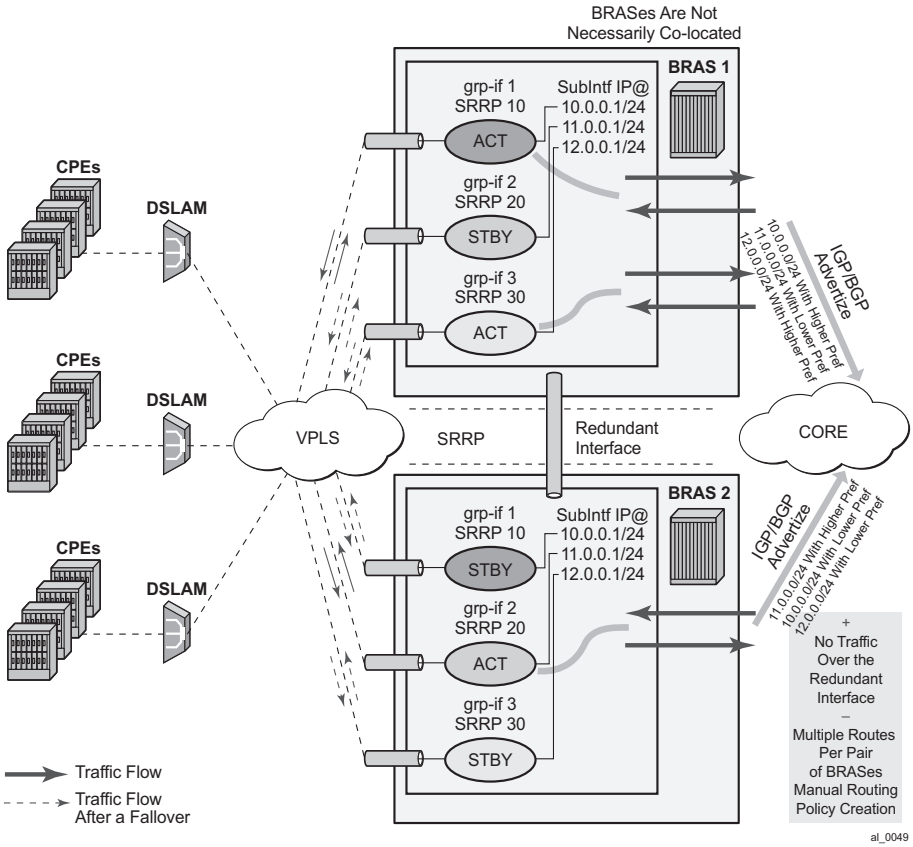

9.3.18.4 MSAP Considerations

As per RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*, this has the implications on the operation of the capture SAP. In an IPoE environment, the initial DHCP traffic related to host establishment uses its native MAC of the physical port on the router. Once the group interface is learned (later in the process, via RADIUS or msap-policy), the MAC address is switched to SRRP MAC address (virtual MAC). The IPoE client adapts easily to this change. On the contrary, for the proper operation of PPPoE with SRRP, the initial destination MAC address learned by the PPPoE client does not change during the lifetime of the session.

This is ensured by indirectly referencing the grp-if under the capture SAP:

```
config>service>vpls
    sap 1/1/1:1.* capture-sap
        track-srrp 10
    sap 1/1/1:2.* capture-sap
        track-srrp 20
```

```
config>service>vprn>
    subscriber-interface <if-name>
        group-interface <grp-if-name>
            sap 1/1/1:1.1
                srrp 10
                message-path 1/1/1:1.1

    group-interface <grp-if-name>
        sap 1/1/1:2.1
            srrp 20
                message-path 1/1/1:2.1
```

With this approach the grp-if is nailed during the session initiation phase by referencing the SRRP instance in track-srrp statement (SRRP is a group interface-wide concept). RADIUS returned grp-if name must match the one on which referenced SRRP instance runs.

The capture SAP of the form

```
sap port-id:*.* capture-sap
    track-srrp X
```

assumes that there is only one grp-if associated with all MSAPs under this capture SAP.

A check is put in place to make sure that the MAC addresses associated with the SRRP instance is the same as the MAC address of the associated capture SAP. A log is raised if there is a discrepancy between the MAC addresses while the grp-if is operationally UP. If there is a MAC address change (user misconfiguration) then the existing PPPoE sessions will time out and the new sessions will fail to establish until the condition is corrected.

9.3.18.5 Unnumbered Interface Support

For unnumbered subscriber-interface support in PPPoE, the gateway IP address that is used to send gratuitous ARP is not available. For this reason, the system IP address is used to send gratuitous ARPs. Gratuitous ARP is used to update the Layer 2 network forwarding path towards the BNG node in the upstream direction.

The system IP address is used automatically if the subscriber interface is unnumbered.

9.3.18.6 Compatibility with MC-LAG

SRRP for PPPoE works in an environment where MC-LAG is enabled. For example, the standby LAG link automatically puts the SRRP node in a Backup state and the SRRP becomes master on the active MC-LAG link. It is important that the SRRP on the standby leg of the MC-LAG is forced into a Backup state, or any new state that will force the downstream traffic to use the redundant interface.

Traffic destined to/from the subscriber is forwarded under the condition that the subscriber-interface is operationally UP. This applies also to shunting of downstream subscriber traffic from the STANDBY to MASTER node. It is always necessary to keep the subscriber-interface operationally UP by configuring a dummy group-interface with a **oper-up-while-empty** command under it. This is especially true for the MC-LAG which causes the messaging SAP on the STANDBY node always to be in the INIT state. In case that MSAPs are used on such group-interfaces, the group-interfaces would be also operationally DOWN, causing the subscriber-interface to be operationally DOWN.

9.3.18.7 IPv6 Support

Prerequisite for MC IPv6 Redundancy is to synchronize PPPoEv6 and IPoEv6 subscribers between the nodes via MCS.

In PPPoE environment, SRRP is used to refresh the forwarding path (MAC addresses) in the access aggregation network (via gratuitous ARP). SRRP ensures that the upstream traffic is steered to the Master BNG node. In the downstream direction, the Backup BNG directs traffic over to the Master BNG node via redundant-interface.

The IPv6 functionality currently relies on IPv4 based SRRP and IPv4 based redundant-interface. In other words, IPv4 is required to run on the access side as well as on the redundant-interface.

The redundant-interface is used in the downstream direction. Traffic arriving on the network links on the Standby node is shunted over to the Master node over the redundant-interface. This is required to ensure consistent QoS and accounting functionality across the nodes (upstream and downstream traffic on the access links for a subscriber must traverse the same BNG node). There is no IPv6 related CLI associated with the redundant-interface.

All IPv6 subscriber traffic that arrives on the Standby node in the downstream direction is automatically shunted over the IPv4 redundant-interface to the Master node. When IPv6 traffic arrives over the redundant-interface on the Master node, it is either PPPoEv6 encapsulated or left as plain IPoEv6 before it is forwarded to the subscriber.

In the upstream direction (AN->BNG) the behavior is the following:

- PPPoEv6

  On the switchover, gratuitous ARPs (gARP) is sent from the new Master on each vlan. The IP address in gARP is the IPv4 gw-ip address or the system IP in the case of unnumbered interfaces. This updates the Layer 2 network path with the proper SRRP MAC address.

- IPoEv6

  IPv4 based SRRP is used to update the Layer 2 forwarding path in the case of a switchover. A gratuitous ARP is sent in the same fashion as it is used for IPoE v4 hosts. Router Advertisements (RA) are not sent out in the case of the switchover.

  However, the two BNG nodes share the same virtual Link Local (LL) IPv6 address. This address is used by the clients as a default-gw and only the Master BNG advertises this LL address in RAs. RAs are suppressed on the Standby node. As already mentioned, RAs are not sent during the switchover. RAs are sent:

  – When the client first gets established – this is how the client learns its default-gw (in PPPoE case RA can also be used for SLAAC – stateless address configuration).

  – As a reply to Router Solicitations messages sent by the clients.

  – Periodically to each client.

  Note that RAs are unicasted to each client.

  Neighbor Advertisements (NA) used for address resolution are sent only from the Master. NA has the SRRP MAC address in the target link layer option on SRRP enabled group interfaces (on non-SRRP enabled group-interfaces, NAs contains the group interface MAC address).

  The LL IPv6 address must be the same on both nodes. In addition, the gw-mac address must be the same on both nodes. The IPv6 clients will not be aware of the switchover and therefore they will not send NS to solicit the update of its neighbor cache with the possibly different gw-mac address.

  The syntax to configure the LL address on the subscriber interface is the following:

```
config>service>ies | vprn>
    subscriber-interface <if-name>
        ipv6
            [no] link-local-address <ipv6-address>
```

```
<ipv6-address> : ipv6-address – x:x:x:x:x:x:x:x:x:x:x:x:d.d.d.d
x [0..FFFF]H
d[0..255]D
```

Note that the current version of SRRP relies only on IPv4 routes. The connection between SRRP and IPv4 routes is done via the subnets with gw IP addresses defined under the subscriber-interfaces in the ESM context. This connection is needed so that SRRP can send Gratuitous ARP properly.

These are the cases for PPPoEv6 MC Redundancy that are supported:

- unnumbered subscriber-interfaces (**config>service>subscriber-interface** hierarchy)
- numbered IPv4 subscriber-interfaces (**config>service>subscriber-interface** hierarchy)
- numbered IPv4 and IPv6 subscriber-interfaces (**config>service>subscriber-interface** and **config>service>sub-if>ipv6** hierarchy)

numbered IPv6 only subscriber-interfaces (**config>service>sub-if>ipv6** hierarchy) is not supported

9.3.18.8 Considerations with Local DHCP Server

When local DHCP Server redundancy/synchronization is used, using address-range failover local | remote, in conjunction with PPPoE in multi-chassis environment, both DHCP servers must be referenced under the corresponding group-interface on each node. For address-range failover access-driven configurations only one DHCP server must be referenced.

```
subscriber-interface <sub-if>
    group-interface <grp-if>
        dhcp
            server <local-dhcp-ip-address> <remote-dhcp-ip-address>
```

Otherwise, the PPPoE clients will not be synchronized via MCS.

Note that this is not the requirement in the IPoE environment. In the IPoE environment, it is enough that the DHCP server points to the IP address of the local DHCP server. If the IP lease is originally assigned by the peer DHCP server, the request for renewal is automatically forwarded to the remote DHCP server by the virtue of the IP address of the original DHCP server that is included in the renewal request.

It is necessary for the successful renewal of the IP address on the remote DHCP server, that the remote DHCP server has a valid return path back to the gi-address of the forwarder of the renewal request.

9.3.18.9 Redundant Interface Considerations

In PPPoE dual-chassis environment without the redundant-interface in place, SRRP aware routing should always be used. Otherwise, if the downstream traffic arrives on the backup node, it will get forwarded directly to the client over the access network (assuming that the access network is operational) with the source MAC address of the group-interface (instead of gw-mac). This grp-if MAC address is different from the MAC address (gw-mac) negotiated during the initial PPPoE phase, and therefore, this traffic will be dropped by the client. It must be ensured that the downstream traffic is always attracted to the Master node in the absence of redundant.

9.3.19 Routed Central Office (CO)

9.3.19.1 Layer 3 Subscriber Interfaces

On regular interfaces in an IES or VPRN service, only one SAP can be associated. A group-interface allows multiple SAPs to be configured as part of a single interface. All SAPs in a single group-interface must be within the same port. Since broadcast is not allowed in this mode, forwarding to the subscriber is based on IP/MAC addresses information gathered by the subscriber management module and stored in the subscriber management table. These entries are based on both static and dynamic DHCP hosts. Routed CO must be used with standard subscriber management or enhanced subscriber management. DSLAMs are typically deployed with Ethernet interfaces.

This model is a combination of two key technologies, subscriber interfaces and group interfaces. While the subscriber interface define the subscriber subnets, the group interfaces are responsible for aggregating the SAPs.

As depicted in Figure 113, an operator can create a new subscriber interface in the IES or VPRN service. A subscriber interface allows for the creation of multiple group interfaces. The IP space is defined by the subnets of the subscriber interface's addresses. Figure 114 shows the details of group interface A.

Figure 113 **Subscriber Interface in an IES/VPRN Service**

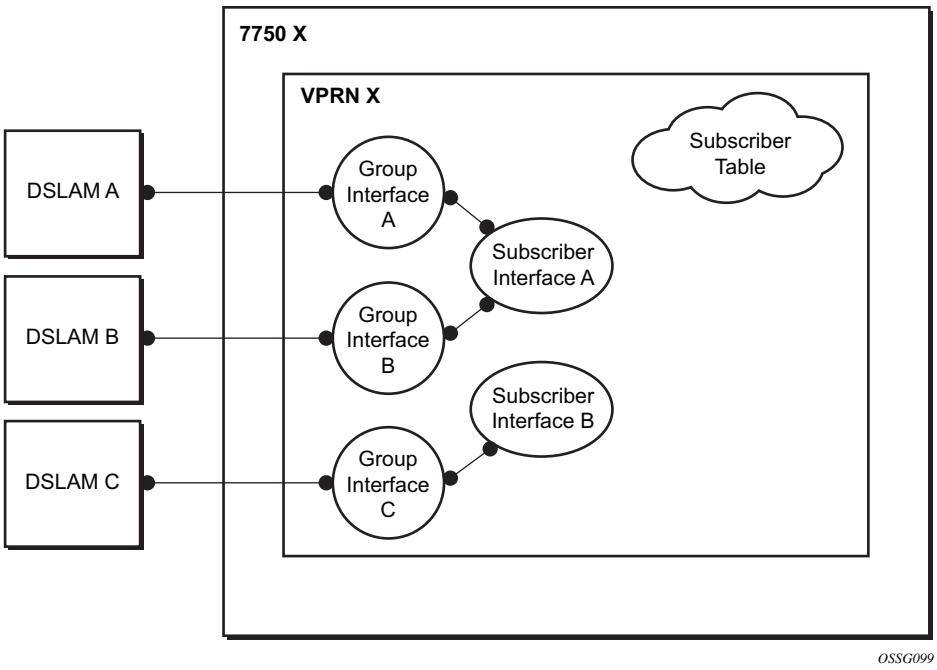

Figure 114 **Details of a Group Interface**

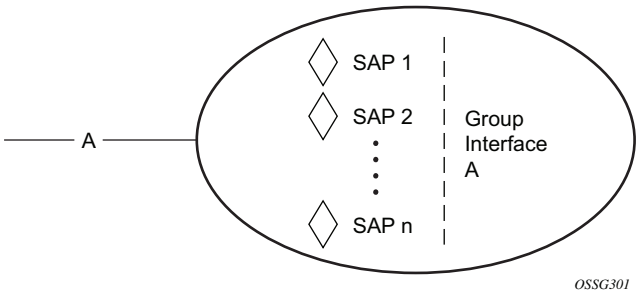

Figure 115 shows a network diagram where the DSLAM are connected directly to a Broadband Service Router (BSR) providing access to an IP subnet. Subscribers from multiple DSLAMs can be part of the same subnet. Note that BSR is also referred to as Broadband Network Gateway (BNG).

Figure 115 **Aggregation Network with Direct DSLAM-BSR Connection**

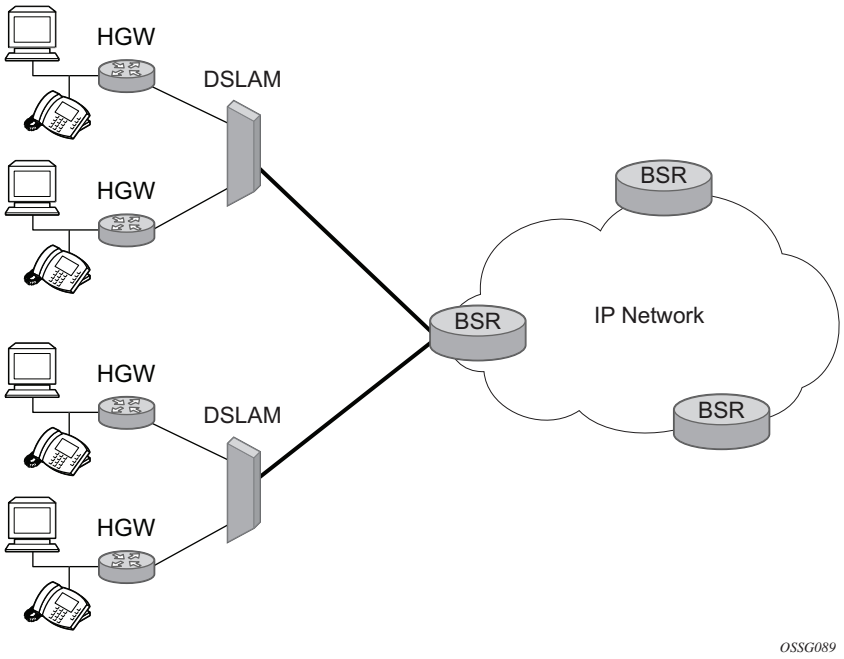

The BSR can be configured with multiple subnets, allowing subscribers to be part of a single subnet as well as providing mechanisms for re-addressing or expanding existing services without affecting existing users.

Figure 116 **Detailed View of Configurable Objects Related to Layer 3 Subscriber Interfaces**

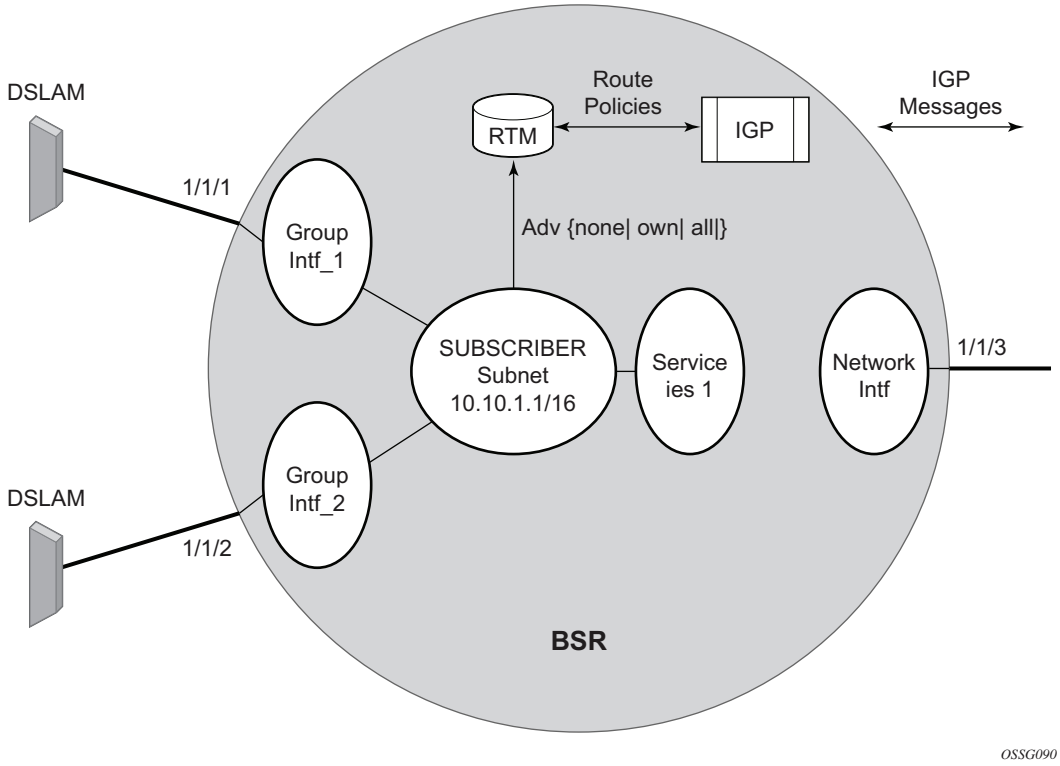

Figure 116 shows a detailed view of a router and the configuration objects implemented to support Layer 3 subscriber interfaces.

- A subscriber service is defined by an IES Service. One or more IES services can be created.
- Each IES service concentrates a number of subscriber-interfaces. The operator can create multiple subscriber interfaces (represented as a subscriber subnet). A subscriber interface will define at least one subnet.
- A group-interface will be provisioned within the subscriber interface for each DSLAM connected. All group interfaces created under the subscriber interface will share the same subnet (or subnets). Group interfaces (shown as intf_1 and intf_2 in Figure 116) are configured as unnumbered and are associated with the subscriber-interface under which they are configured.
- SAPs can be configured under the group-interface. In a VLAN-per-DSLAM model only, one SAP per group-interface is needed, while in the VLAN-per-subscriber model, a subscriber of the DSLAM will require its own SAP. All SAPs on a group-interface must be on the same physical port or LAG.

The individual features related to subscribers, such as DHCP relay, DHCP snooping and anti-spoofing filters, are enabled at group-interface level. For a Routed CO model of subscriber management, and when enhanced subscriber management (if sub-sla-mgmt is configured). Then, hashing will be based on an internally assigned subscriber-ID. Having a unique subscriber ID configured in CLI will ensure that each subscriber is assigned a unique internal subscriber ID.

It is assumed that individual end-user devices (further referred to as subscriber hosts) get their IP address assigned through either DHCP or static configuration. The management of individual subscriber hosts (such as creation, queue allocation, and so on) is performed by Enhanced Subscriber Management.

The operator can provision how the system advertises routes. While most deployments will advertise the full subnet it is possible to have the system advertise only the active, discovered or static host routes.

The distribution of this information into routing protocols will be driven by import/ export route policies configured by the operator.

9.3.19.1.1 DHCP Interactions

The DHCP relay process has been enhanced to record incoming DHCP discover and request messages. Since forwarding to the SAPs is done by the information in the subscriber management table and multiple SAPs are allowed in one interface it was impossible to know which SAP will be used to forward the DHCP replies. The node maintains a cache of the DHCP requests. The cache can be viewed using the **tools>dump>router>dhcp>group-if-mapping** command. The cache holds an entry for 30 seconds. If an ACK/NAK packet was not received from the server within the timeout the node discards the cache entry. The node can use the Option 82 circuit-id field as part of the temporary host entry. If used, the ACK must contain the same circuit-id field in Option 82 to be found in the cache only if the match-circuit-id is specified at the DHCP level of the group- interface. When the **match-circuit-id** command is enabled a check is performed for option 82 circuit-id.

9.3.19.1.2 Routed CO for IES Service

The routed CO model depends on subscriber management to maintain the subscriber host information. To create a group-interface the operator must first create a subscriber interface within the service (**config>service>ies>subscriber-interface** *ip-int-name*). The subscriber interface maintains up to 256 subscriber subnets and is configured with a host address for each subnet.

When a DHCP ACK is received the IP address provided to the client will be verified to be in one of the subscriber subnets associated with the egress SAP. It will be noted that when DHCP snooping is enabled for regular IES interfaces the same rule will apply.

The subscriber interface is an internal loopback interface. The operational state is driven from the child's group-interface states and the configuration of an address in the RTM.

The group interface is an unnumbered interface. The interface will be operationally up if it is in the no shutdown state and if at least one SAP has been defined and is up and the parent subscriber interface is administratively up. The first SAP defined will determine the port for the group- interface. If the user attempts to define a subsequent SAP that is on a different port will result in an error. When the subscriber-interface or the group-interface is in shutdown state no packets will be delivered/ received to/from the subscriber hosts but the subscriber hosts, both dynamic and static, will be maintained based on the lease time.

In the routed CO model, the router acts as a DHCP relay agent and also serves as the subscriber- identification agent. The DHCP actions are defined in the group-interface. All SAPs in that interface inherit these definitions. The group-interface DHCP definition are a template for all SAPs.

Lease-populate is enabled by default with the number-of-entries set to 1. This enables DHCP lease state for each SAP in the group-interface.

Since the group-interface can aggregate subscribers in different subnets a GI address must be defined for the DHCP relay process. The address must be in one of the host addresses defined for the subscriber interface. The GI address can be defined at the subscriber interface level which will cause all child group interface to inherit that route. The GI address can then be overridden at the group interface level. A GI address must be defined in order for DHCP relay to function.

Because of the nature of the group-interface, local-proxy-arp, as well as arp-populate, should be enabled. This would allow the system to respond to subscriber ARP requests if the ARP request contains an IP address which is in the same subnet as one of the subscriber interface subnets.

When an authentication policy is specified for a SAP under a group interface, DHCP will intercept DHCP discover messages for RADIUS authentication. If the system is a DHCP-relay defined in a group-interface and the GI address was not configured the operational state of DHCP will be down.

9.3.19.1.3 Routed CO for VPRN Service

Much like in Routed CO for IES service, the Routed CO model for VPRN depends on subscriber management to maintain the subscriber host information. To create a group-interface, the operator must first create a subscriber interface in the config>service>vprn context. The subscriber interface can maintain up to 256 subscriber subnets and can be configured with a host address for each subnet. The host IP address can be installed as a result of both relaying to a DHCP server and proxy to a RADIUS server. In both cases the host IP address must be in the subnet defined by the VPRN's subscriber interface.

Basic subscriber management is allowed only in a subscriber/SAP model and can be used only in a dedicated VPRN architecture. A RADIUS service selection (using Managed SAPs) will require Enhanced Subscriber Management. The subscriber interface's subnets are allowed to be advertised to both IGPs and BGP within a VPRN.

When an authentication policy is specified for a group-interface, DHCP snooping must be enabled to intercept DHCP discover and renew messages for RADIUS authentication. Subscriber management RADIUS extensions are allowed if the operator chooses to have the RADIUS server return the subscriber identification, subscriber profile and sla-profile strings using RADIUS.

The node can be defined with both a DHCP relay or proxy function. If the user configures a DHCP relay, the **local-proxy-server** command will enable DHCP split leases. In that configuration the node will provide the configured DHCP lease to the client using either RADIUS or the real DHCP server as the source of the IP address to be provided.

The RADIUS server can send a Change of Authorization (CoA) message containing the DHCP FORCERENEW VSA which prompts the local-proxy-server to send a FORCERENEW message to the client. The node ACKs when the FORCERENEW messages has been sent, regardless of whether the subscriber responds. If the client fails to respond or if a new session cannot be established due to resource management issues or otherwise the node must respond with a NACK to the RADIUS server.

If the CoA message contains an IP address that is different than the configured IP address (when RADIUS was providing IP addresses) the node must send a FORCERENEW message to the client and NAK the request and provide a new IP address. If the node fails to receive a request, the CoA is ACK'd when the FORCERENEW message has been sent.

The operational state of group and subscriber interfaces are dependent on the state of active SAPs. A group interface can become operationally up only if at least one SAP is configured and is in an operationally up state. A subscriber interface becomes operationally up if at least one group interface is operationally up or the associated wholesale forwarding interface is operationally up. This ensures that, in a failure scenario that affects all group interfaces in a given subscriber subnet, the node will stop advertising the subnet to the network. The SRRP state will affect this behavior as well and can cause the subnet to be removed if all group interfaces (and SRRP instances) are in backup state.

9.3.19.2 Wholesale Retail Routed CO

VPRN Routed CO allows a provider to resell wholesaler services (from a carrier) while providing direct DSLAM connectivity. An operator can create a VPRN service for the retailer and configure the access from subscribers as well as to the retailer network. Any further action will be as if the VPRN is a standalone router running the Routed CO model. All forwarding to these servers must be done within the VPRN service. The operator can leak routes from the base routing instance. In this model, the operator can use RADIUS for subscriber host authentication, DHCP relay and DHCP proxy. This provides maximum flexibility to the retailer while minimizing the involvement of the wholesaler. Access cannot be shared among retailers unless a subscriber SAP is used. This requires that the wholesaler maintain a different access node (DSLAM) for each retailer that does not scale well. The wholesale retail model described in this section overcomes these limitations.

9.3.19.2.1 Wholesale Retail Model

In the wholesale retail model (Figure 117), the wholesaler instance connections that are common to the access nodes are distributed to many retail instances. A subscriber host attached to an access node connected in the wholesaler service can be instantiated in a retail service and obtain IP addresses from the retailers address space. The service context of the retailer is determined during the subscriber host authentication phase (for example via the Alc-Retail-Serv-Id attribute in RADIUS or the retail-service-id CLI in the local user database).

Upstream subscriber traffic ingresses into the wholesaler instance and after identification is then forwarded into the retail instance. The reverse will occur for traffic in the downstream direction.

Figure 117 **Wholesale Retail Model**

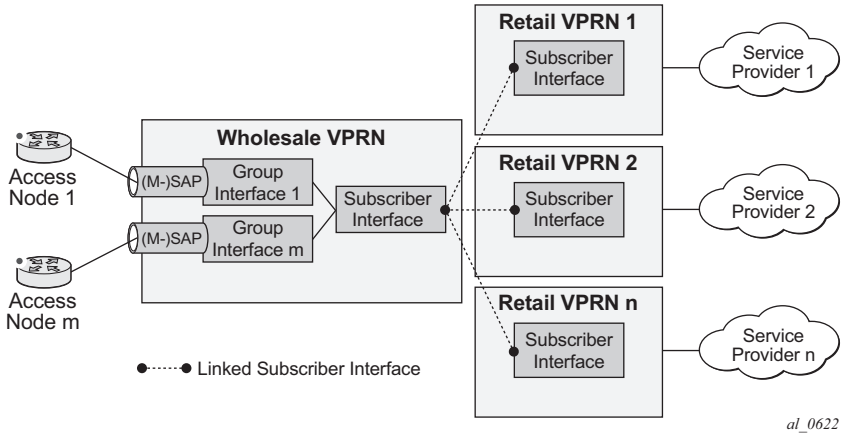

In a wholesale retail model, two subscriber interfaces must be configured and linked together: one in the wholesale VPRN and one in the retail service.

The wholesale subscriber interface defines the IP subnets and host specific configuration parameters for subscriber hosts belonging to the wholesaler. There are associated group interfaces that contain the SAPs which connect to the access nodes.

The retail subscriber interface defines the IP subnets and host specific configuration parameters for subscriber hosts belonging to the retailer. The retail subscriber interface is linked to a wholesale subscriber interface for forwarding via explicit configuration. There are no associated group interfaces.

For example:

```
config>service
        vprn 1000 customer 1 create
            subscriber-interface "sub-int-ws-1" create
            # wholesale subscriber interface
                --- snip ---
                group-interface "group-int-1-1" create
                    --- snip ---
                    sap 1/1/1:1 create
                        --- snip ---
                    exit
                exit
            exit
        exit


       vprn 1001 customer 1 create
            subscriber-interface "sub-int-rt-1" fwd-service 1000    \\
                                 fwd-subscriber-interface "sub-int-ws-1" create
            # linked retail subscriber interface
                --- snip ---
```

```
        exit
    exit
```

A retail subscriber interface can be linked to a single wholesale subscriber interface and context only. Subscriber interface chaining (linking a retail subscriber interface to another retail subscriber interface) is not supported. Multiple retail subscriber interfaces belonging to different retail contexts can be associated with a single wholesale subscriber interface. When a retail subscriber interface is linked to a wholesale context, all other retail subscriber interfaces from the same retailer must be linked to the same wholesale context.

9.3.19.2.2 Configuration and Applicability

As explained in the previous section, the wholesale retail model is provisioned with the linking of a subscriber interface in a retail service to a subscriber interface in the wholesale VPRN service.

Because a retail subscriber interface does not have a group-interface context, some group-interface-specific CLI parameters such as to **configure dhcp relay** are made available at the retail subscriber interface level. Other CLI parameters such as to provision Radius or local user database authentication are configured at the wholesale subscriber or group interface and apply to both wholesale and retail subscriber hosts.

The DHCP lease-populate configuration is special in wholesale retail as it is configured in both wholesale and retail context. The lease-populate value in the wholesale group-interface dhcp context controls the per SAP limits while the lease-populate value configured in the retail subscriber interface dhcp context controls the limits for the retailer subscriber interface. Both limits must be satisfied before a new subscriber host can be instantiated.

The sample configurations below enable dual stack IPoE devices to connect to wholesale service VPRN 4000 and retail service VPRN 4001. Hosts connected in VPRN 4000 get their IP address assigned from RADIUS, hence the proxy server configuration. Hosts connected in VPRN 4001 get their IP address from a DHCP server, hence the DHCP relay configuration.

Only the service configurations are shown. They have to be completed with authentication policies and subscriber management configuration such as radius-server-policies, sub- and sla-profiles, and so on.

Sample configuration – Wholesale VPRN Service:

```
config>service
```

```
vprn 4000 customer 1 create
    autonomous-system 64500
    route-distinguisher 64500:4000
    auto-bind-tunnel
        resolution-filter
            ldp
            rsvp
        exit
        resolution filter
    exit
    vrf-target target:64500:4000
    subscriber-interface "sub-int-1" create
        address 10.10.1.254/24
        address 10.10.2.254/24
        ipv6
            delegated-prefix-len variable
            subscriber-prefixes
                prefix 2001:db8:a:100::/56 wan-host
                prefix 2001:db8:a001::/48 pd
            exit
        exit
        group-interface "group-int-1" create
            ipv6
                router-advertisements
                    no shutdown
                exit
                dhcp6
                    proxy-server
                        no shutdown
                    exit
                exit
            exit
            arp-populate
            dhcp
                proxy-server
                    emulated-server 10.10.1.254
                    no shutdown
                exit
                lease-populate 100
                no shutdown
            exit
            authentication-policy "auth-policy-1"
            sap 1/1/4:1201.27 create
                sub-sla-mgmt
                    sub-ident-policy "sub-ident-1"
                    multi-sub-sap 100
                    no shutdown
                exit
            exit
        exit
    exit
    no shutdown
exit
```

Sample configuration – Retail VPRN Service:

```
config>service>
        vprn 4001 customer 1 create
            autonomous-system 64501
            route-distinguisher 64500:4001
            auto-bind-tunnel
                resolution-filter
                    ldp
                    rsvp
                exit
                resolution filter
            exit
            vrf-target target:64500:4001
            interface "int-loopback-1" create
                address 192.0.2.5/32
                ipv6
                    address 2001:db8::5/128
                exit
                loopback
            exit
            subscriber-interface "sub-int-rt-4000-1" fwd-service 4000 fwd-
subscriber-
interface "sub-int-1" create
                address 10.10.11.254/24
                address 10.10.12.254/24
                dhcp
                    server 192.0.2.4
                    lease-populate 100
                    gi-address 10.10.11.254
                    no shutdown
                exit
                ipv6
                    subscriber-prefixes
                        prefix 2001:db8:b:100::/56 wan-host
                        prefix 2001:db8:b001::/48 pd
                    exit
                    dhcp6
                        relay
                            source-address 2001:db8::5
                            server 2001:db8::4
                            no shutdown
                        exit
                    exit
                    router-advertisements
                        no shutdown
                    exit
                exit
            exit
            no shutdown
        exit
```

The wholesale retail model applies to all IPoE, PPPoE PTA, IPv4 and IPv6 host types.

The wholesale service type must be VPRN. For IPoEv4 hosts, the retail service type must be a VPRN. For all other host types, the retail service type can be IES or VPRN.

Multicast-per-host replication can be enabled without support for multi-chassis redundancy.

The wholesale retail model can be deployed in combination with managed SAPs.

Overlapping subscriber subnets and prefixes in retail VPRN services associated with the same wholesale forwarding service are supported for PPPoE (IPv4 and IPv6) and IPoE (IPv4 and IPv6). This support is enabled by configuring private retail subnets on the retail subscriber interface. Private retail subnets cannot be configured when multi-chassis redundancy is needed.

9.3.19.2.3 Hub-and-Spoke Forwarding

In some cases, hub-and-spoke-type forwarding is needed for the retailer's VPRN. When the retailer expects all subscriber traffic to reach its router (for accounting, monitoring, wiretapping, and so on) normal best-hop behavior within the retailer VPRN is not desired. Any subscriber-to-subscriber traffic is forwarded within the VPRN preventing the retailer from receiving these packets. To force all subscriber packets to the retailer network, a new type of hub-and-spoke topology is defined: "type subscriber-split-horizon". It can be used to force all subscriber traffic (upstream) to the retailer's network. The system requires that the operator shutdown the VPRN service to enable this flag.

With retail VPRN type configured to subscriber-split-horizon, routes learned from MBGP, IGP through a regular interface, static routes through regular interfaces and locally attached regular interface routes are considered hub routes and are used for upstream traffic forwarding. Subscriber subnets cannot be used for upstream traffic forwarding. Downstream traffic will use routes in both hub and spoke routing instances.

Figure 118 shows user-to-user traffic forwarding for both retail VPRN types: regular and subscriber-split-horizon.

Figure 118 **Wholesale Retail – Hub-and-Spoke Forwarding**

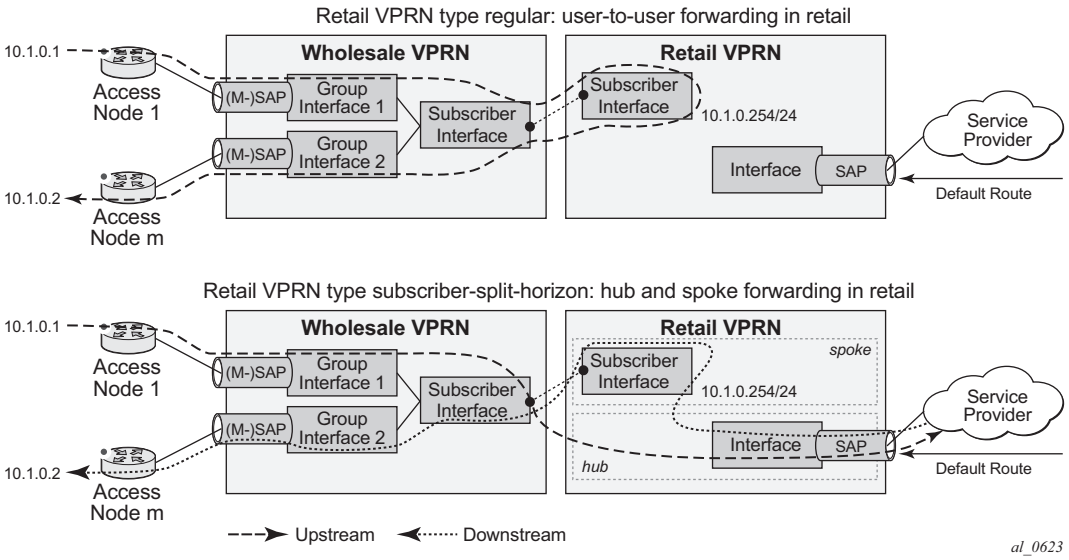

Hub-and-spoke forwarding can also be used in combination with wholesale unicast RPF (uRPF) check. The uRPF is performed on upstream traffic on spoke routes (subscriber subnets) and the forwarding uses hub routes only.

9.3.19.3 Routed Subscriber Hosts

Figure 119 **Router Subscriber Hosts**

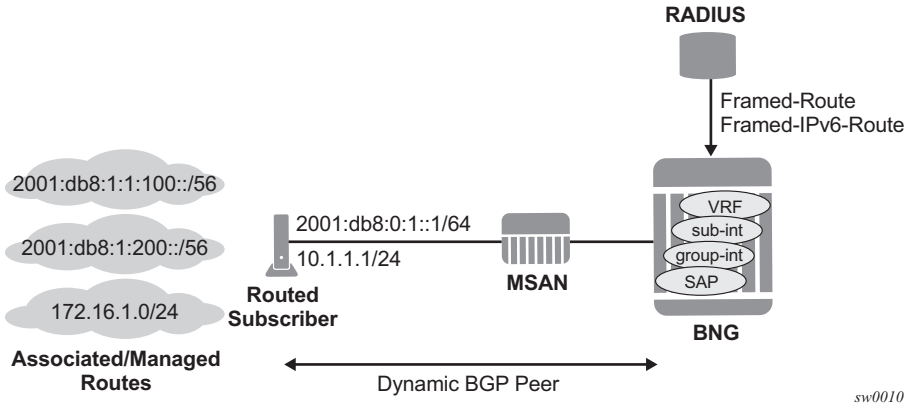

A routed subscriber host associated route is a global routable subnet/prefix behind a routed CPE or Home Gateway. The routed CPE is identified in the BNG as an ESM subscriber host: QoS, accounting and anti-spoofing is enforced per CPE. The associated routes are installed in the BNG route table with next-hop pointing to the routed subscriber host's WAN address.

Routed subscriber host associated routes are supported on IES/VPRN subscriber interfaces in a routed CO configuration. To put a SAP or MSAP in routed subscriber mode, the anti-spoof type for the SAP or MSAP must be configured to nh-mac:

```
configure
    service ies/vprn <service-id>
        subscriber-interface <ip-int-name>
            group-interface <ip-int-name>
                sap <sap-id>
                           anti-spoof nh-mac

configure
    subscriber-mgmt
        msap-policy <msap-policy-name>
            ies-vprn-only-sap-parameters
                anti-spoof nh-mac
```

There are three ways to learn about a routed subscriber host associated IPv4 route:

1. Configuration for a static host
2. A dynamic BGP peer
3. The RADIUS [22] Framed-Route attribute

A routed subscriber host associated IPv6 route can only be learned with the RADIUS [99] Framed-IPv6-Route attribute.

9.3.19.3.1 Static Configured IPv4 Managed Route

The routes associated with a static host are populated in the routing table as "Remote Managed" routes. Up to sixteen managed routes can be configured for a static host.

```
config>service>ies>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>sap

    static-host ip 10.1.1.20 create
        sla-profile "sla-profile-1"
        sub-profile "sub-profile-1"
        subscriber "static-host-1"
        managed-routes
            route 172.20.1.0/24
```

```
            . . .
            route 172.20.16.0/24
        exit
        no shutdown
           exit
```

To display the managed routes associated with a routed subscriber host, use following commands:

show service id *service-id* **static-host detail**

9.3.19.3.2 Static Configured IPv6 Managed Route

The routes associated with a static host are populated in the routing table as “Remote Managed” routes. Up to sixteen managed routes can be configured for a static host.

```
config>service>ies>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>sap
    anti-spoof nh-mac
    static-host ip 2001::1/128 create
        sla-profile "sla-profile-1"
        sub-profile "sub-profile-1"
        subscriber "static-host-1"
        managed-routes
            route 2000::/56
. . .
            route 3000::/56
        exit
    no shutdown
    exit

    static-host ip 2001:1::/64 create
    sla-profile "sla-profile-1"
    sub-profile "sub-profile-1"
    subscriber "static-host-1"
    managed-routes
        route 4000::/56
. . .
        route 5000::/56
        exit
    no shutdown
    exit
```

To display the managed routes associated with a routed subscriber host, use following commands:

show service id *service-id* **static-host detail**

9.3.19.3.3 Dynamic BGP Peering

Routed subscriber host associated IPv4 routes can be learned over a dynamic BGP IPv4 peer that is automatically set up when a subscriber host is instantiated. The parameters for the BGP peer are configured in a BGP peering policy or obtained in Radius VSA attributes. The subscriber-host IPv4 address is used as the BGP peer IP address. The BGP peering is torn down and the associated routes removed from the routing table as soon as the subscriber-host is removed.

Dynamic BGP peering is supported for routed subscriber hosts terminated in a VPRN service and is not supported for routed subscriber hosts terminated in an IES service. The BGP learned routes scaling is limited by the BGP scaling limits. The routes learned via a dynamic BGP peer are populated in the routing table as “Remote BGP” routes.

To display the BGP learned routes associated with a routed subscriber host, use the regular BGP commands. For example:

show router *service-id* **bgp neighbor** *ip-address* **received-routes**

An ESM dynamic BGP group must be configured in the BGP cli context of the VPRN service where the subscriber host is started:

config>service>vprn>bgp

```
    group "dynamic-peer-1" esm-dynamic-peer
    exit
```

The BGP peering policy to be used must be configured in the subscriber-mgmt CLI context:

```
config>subscr-mgmt

    bgp-peering-policy "bgp-policy-1" create
    exit
```

An ESM dynamic BGP peer is established for a subscriber host if the RADIUS attribute [26-6527-55] “Alc-BGP-Policy” returned in the Access-Accept contains the name of a local configured bgp-peering-policy and if an ESM dynamic peer group is configured in the VPRN BGP context.

BGP peering parameters can be specified from multiple sources:

- Use BGP peering parameters returned in RADIUS VSA attributes
- If not available from RADIUS, use BGP peering parameters configured in the bgp-peering-policy

- If not configured in the bgp-peering-policy, use BGP peering parameters configured for the esm-dynamic-peer group
- If not configured in the esm-dynamic-peer group, use the BGP peering parameters configured in the VPRN service BGP CLI context.
- If not configured in the VPRN service BGP CLI context, use the defaults

The import and export policies to be used for the dynamic bgp peer are determined in following priority order:

1. Use import/export policies returned in RADIUS VSA attributes and append policies configured in the bgp-peering-policy.
2. If not available from RADIUS and not configured in the bgp-peering-policy, use the policies configured in the esm-dynamic-peer group.
3. If not configured in the esm-dynamic-peer group, use the policies configured in the VPRN service BGP CLI context.

Table 83 details the RADIUS VSA attributes that can be used to setup dynamic BGP peering.

Table 83 **RADIUS VSA Attributes to Setup Dynamic BGP Peering**

Attribute-ID	Attribute Name	Description
26-6527-55	Alc-BGP-Policy	Mandatory attribute to setup a dynamic BGP peer. References a bgp peering policy configured in the “configure subscriber-mgmt bgp-peering-policy <policy-name>” CLI context.
26-6527-56	Alc-BGP-Auth-Keychain	Optional. References a keychain configured in the “configure system security keychain <keychain-name> “ CLI context.
26-6527-57	Alc-BGP-Auth-Key	Optional. The MD5 authentication key used between BGP peers for BGP session establishment.
26-6527-58	Alc-BGP-Export-Policy	Optional. References a pre-configured BGP export routing policy.
26-6527-59	Alc-BGP-Import-Policy	Optional. References a pre-configured BGP import routing policy.
26-6527-60	Alc-BGP-PeerAS	Optional. Specifies the Autonomous System number for the remote peer

9.3.19.3.4 RIP Listener

If a routed subscriber host is associated with a RIP policy, the host's IPv4 routes can be learned over RIP. The BNG only supports RIP listener and does not support sending RIP routes to subscribers. To enable RIP for a subscriber, the subscriber must first be associated with a **rip-policy**. The group interface of the subscriber must also be configured as a RIP neighbor. The rip- policy can be associated to the subscriber during authentication from LUDB or via RADIUS. It can also be configured directly for static hosts. The RIP routes learned from a subscriber is removed as a subscriber is purged or shutdown from the system. RIP listening for ESM host is supported on both IES and VPRN.

To display the RIP learned routes associated with a routed subscriber host, use the regular RIP commands. For example:

```
show router service-id rip neighbor interface advertised-routes
```

The group interface must be configured in the RIP CLI context of the routed instance where the subscriber host is created:

```
config>router/service vprn>rip
    group "rip-listener"
        neighbor "group-interface-01"
```

The RIP policy is configured in the subscriber-mgmt CLI context:

```
config>sub-mgmt
    rip-policy "rip-policy-01" create
```

A RIP neighbor is established for a subscriber host if the RADIUS attribute [26-6527-207] "Alc-RIP-Policy" is returned in the Access-Accept or in LUDB. RIP parameters such as **authentication key** and **type** can be specified in the RIP policy.

For more information about RIP, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Unicast Routing Protocols Guide*.

9.3.19.3.5 RADIUS: Framed-Route and Framed-IPv6-Route

RADIUS attribute [22] Framed-Route can be specified in a RADIUS Access-Accept message to associate an IPv4 route with an IPv4 routed subscriber host and Radius attribute [99] Framed-IPv6-Route can be used to associate an IPv6 route with an IPv6 routed subscriber wan host (DHCPv6 IA-NA or SLAAC). These routes are populated in the routing table as “Remote Managed” routes. Up to sixteen managed routes can be installed for a routed subscriber host; this corresponds with up to sixteen Framed-Routes and sixteen Framed-IPv6-Routes for a dual stack routed subscriber. Framed-IPv6-Routes cannot be associated with a Prefix Delegation host (DHCP IA-PD).

The Framed-Route and Framed-IPv6-Route attributes should be formatted as:

"<ip-prefix>[/<prefix-length>] <space> <gateway-address> [<space> <metric>] [<space> tag <space> <tag-value>] [<space> pref <space> <preference-value>]”

where:

<space> — is a white space or blank character.

<ip-prefix>[/prefix-length] — is the managed route to be associated with the routed subscriber host. The prefix-length is optional for an IPv4 managed route. When not specified, a class-full class A,B or C subnet is assumed. The prefix-length is mandatory for an IPv6 managed route.

<gateway-address> — must be the routed subscriber host IP address. “0.0.0.0” is automatically interpreted as the host IPv4 address for managed IPv4 routes.

“::” and “0:0:0:0:0:0:0:0” are automatically interpreted as the wan-host IPv6 address for managed IPv6 routes.

[<metric>] — Optional. Installed in the routing table as the metric of the managed route. If not specified, metric zero is used. Value = [0.. 65535].

[tag <tag-value>] — Optional. The managed route will be tagged for use in routing policies. If not specified, or tag-value = 0, then the route is not tagged. Value = [0..4294967295].

[pref <preference-value>] — Optional. Installed in the routing table as protocol preference for this managed route. If not specified, preference zero is used. Value = [0..255].

If the optional metrics (metric, tag and/or preference) are specified in a wrong format or with out of range values, then the defaults are used for all metrics: metric=0, no tag and preference=0. No event is logged.

If the Framed-Route or Framed-IPv6-Route is invalid (for example because the gateway address specified does not match the host wan IP address or because the host bits are not zero) then the routed subscriber host is instantiated without the ill-defined managed route. An event is logged in this case.

Equal Cost Multi-Path (ECMP) is supported for Framed-Route and Framed-IPv6-Route:

The maximum number of equal cost paths in a routing instance is configured with:

```
config>router>
config>service>vprn>
        ecmp <max-ecmp-routes>
```

If an identical managed route is associated with different routed subscriber hosts in the context of the same IES/VPRN service, up to <*max-ecmp-routes*> managed routes are installed in the routing table. Candidate ECMP Framed-Routes/Framed-IPv6-Routes have:

- Identical prefix
- Equal lowest preference
- Equal lowest metric

A tie breaker determines if more candidate ECMP Framed-Routes/Framed-IPv6-Routes are available than the configured <max-ecmp-routes> is: Lowest ip next-hop.

Other identical managed routes are shadowed and an event is logged.

Note that Candidate ECMP Framed-Routes/Framed-IPv6-Routes can belong to hosts of the same or different subscriber.

Valid Framed-Routes and Framed-IPv6-Routes are persistent (stored in the persistency file for recovery after reboot) and synchronized in a Multi-Chassis Redundancy configuration.

RADIUS-learned Framed-Route/Framed-IPv6-Route and static host associated managed routes that are installed in the routing table can be identified in routing policies for redistribution as protocol “managed”.

To display the managed routes associated with a routed subscriber host, use following commands:

show service id *service-id* **dhcp lease-state detail**

show service id *service-id* **dhcp6 lease-state detail**

show service id *service-id* **slaac host detail**

show service id *service-id* **ppp session detail**

show service id *service-id* **pppoe session detail**

show service id *service-id* **arp-host detail**

Valid RADIUS-learned managed routes can be included in RADIUS accounting messages with the following configuration:

```
configure
    subscriber-mgmt
        radius-accounting-policy <name>
            include-radius-attribute
                framed-route
                framed-ipv6-route
```

Associated managed routes for an instantiated routed subscriber host are included in RADIUS accounting messages independent of the state of the managed route (Installed, Shadowed, HostInactive, and so on)

For a PPP session, when a Framed-Route or Framed-IPv6-Route is available while the corresponding routed subscriber host is not yet instantiated, the managed route is in the state “notYetInstalled” and will not be included in RADIUS accounting messages.

9.3.19.3.6 GRT Lookup and Routed CO in a VPRN

GRT lookup allows routing from a VPRN to the GRT, and GRT leaking allows routing from the GRT to a VPRN. These features are particularly useful when VPRNs require routing to the Internet and the GRT already contains the Internet routing table. Wholesale/retail VPRNs and the routed CO VPRN have both GRT lookup and GRT leaking support.

The **config>service>vprn>grt-lookup>export-grt** command exports subscriber-related routes and protocols to the GRT. This allows traffic arriving from the network port to be routed downstream to the subscriber. The following configurations are supported in the downstream direction.

- For an IPv4 numbered subscriber interface inside a routed CO VPRN, an IPv4 subscriber subnet can be exported as a policy using the **config>service>vprn>grp-lookup>export-grt** command, where the policy is configured with the **config>router>policy-options>policy-statement** command.

- For an IPv6 numbered subscriber interface inside a routed CO VPRN, an IPv6 subscriber subnet/prefix can be exported as a policy using the **config>service>vprn>grp-lookup>export-grt** command, where the policy is configured with the **config>router>policy-options>policy-statement** command.
- Subscriber-related protocols (managed routes and subscriber management routes) inside a routed CO VPRN can be exported to the GRT.
- For an IPv4 unnumbered subscriber interface inside a routed CO VPRN, subscriber host /32 routes are exported individually to the GRT using the “sub-mgmt” protocol.
- For an IPv6 unnumbered subscriber interface inside a routed CO VPRN, subscriber host /128 routes and/or prefixes are exported individually to the GRT using the “sub-mgmt” protocol.
- The retail VPRN supports all items listed above for a routed CO VPRN (exporting subscriber routes and protocols to the GRT via the **config>service>vprn>grp-lookup>export-grt** command). The retail VPRN does not export host routes by default. Therefore, the **export-host-routes** command may be required for a retail VPRN unnumbered subscriber interface.
- The wholesale VPRN supports all items listed above for a routed CO VPRN (exporting subscriber routes and protocols to the GRT via a route policy). However, retail-related routes that appear in the FDB of the wholesale VPRNs, cannot be exported. Retail-related routes must be exported individually within the VPRN to the GRT. This provides control over which retail VPRNs route to the GRT.

GRT lookup supports traffic from the subscriber to be routed upstream to the GRT. The following configurations are supported in an upstream direction:

- routed CO VPRN and **grt-lookup**
- wholesale VPRN and **grt-lookup**
- retail VPRN and **grt-lookup**

Not Supported

- SRRP setup in routed CO VPRN
- SRRP setup in wholesale/retail VPRN

9.3.20 Dual Homing

All residential networks are based on two models: Layer 2 CO and Layer 3 CO. Dual homing methods for Layer 2 CO include MC-LAG and MC-Ring. Dual homing for Layer 3 CO is based on SRRP and can be done in ring-topologies (l3-mc-ring or with directly attached nodes. All methods use multi-chassis synchronization protocol to sync subscriber state.

9.3.20.1 Dual Homing to Two PEs (Redundant-Pair Nodes) in Triple Play Aggregation

Figure 120 **Dual-Homing to Two PEs**

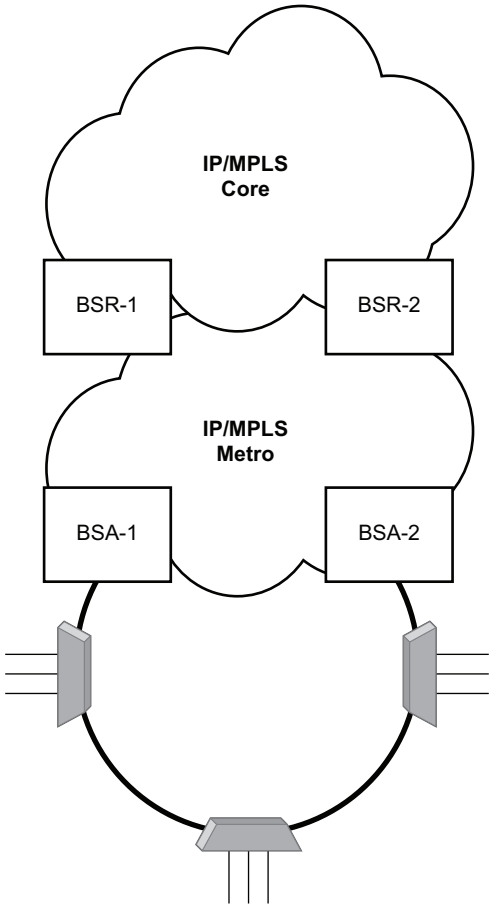

Figure 120 depicts dual-homing to two different PE nodes. The actual architecture can be based on a single DSLAM having two connections to two different PEs (using MC-LAG) or ring of DSLAMs dual-connected to redundant pair of PEs.

Similarly to previous configuration, both aggregation models (VLAN-per-subscriber or VLAN-per-service) are applicable.

Configurations include:

- Loop resolution and failure recovery — Can be based on MC-LAG or mVPLS.
- DHCP-lease-state persistency — Stores all required information to recover from node failure.
- DHCP-lease-state synchronization — A mechanism to synchronize the DHCP lease-state between two PE nodes in the scope of redundancy groups (a group of SAPs used for dual homing).
- IGMP snooping state synchronization — Similarly to DHCP lease-state synchronization, IGMP snooping state is synchronized to ensure fast switchover between PE nodes. In a VPLS network, a BTV stream is typically available in all PE nodes (the ring interconnecting all PEs with Mcast routers is typically used) so the switch over can be purely driven by RSTP or MC-LAG.
- ARP reply agent responses — The ARP reply agent can response to ARP requests addressing a host behind the given SAP if the SAP is in a forwarding state. This prevents the FDB table in the VPLS from being “poisoned” by ARP responses generated by the node with a SAP in a blocking state (see Figure 121).

Figure 121 shows a typical configuration of network model based on Layer 2 CO model. Individual rings of access nodes are aggregated at BSA level in one (or multiple) VPLS services. At higher aggregation levels (the BSR), individual BSAs are connected to Layer 3 interfaces (IES or VPRN) by spoke SDP termination. Every Layer 3 interface at BSR level aggregates all subscribers in one subnet.

Figure 121 **Layer 2 CO Dual Homing - Network Diagram**

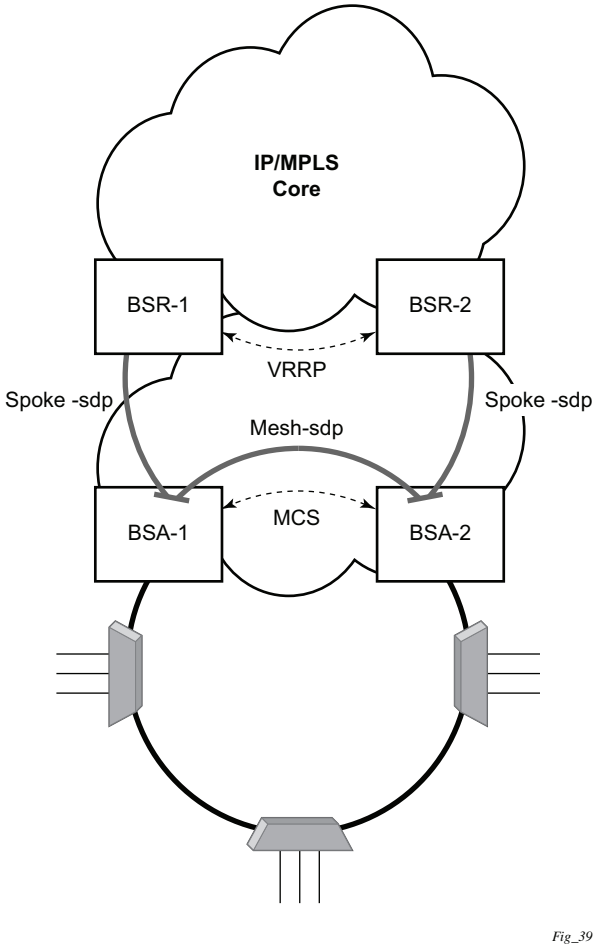

Typically, BTV service distribution is implemented in a separate VPLS service with a separate SAP per access-node. This extra VPLS is not explicitly indicated in Figure 121 (and subsequent figures) but the descriptions refer to its presence.

From a configuration point of view in this model, it is assumed that all subscriber management features are enabled at the BSA level and that synchronization of the information (using multi-chassis synchronization) is configured between redundant pair nodes (BSA-1 and BSA-2 shown in Layer 2 CO Dual Homing - Network Diagram). The multi-chassis synchronization connection is used only for synchronizing active subscriber host database and operates independently from dual-homing connectivity control. At the BSR level, there are no subscriber management features enabled.

The operation of redundancy at the BSR level through VRRP is the same as dual homing based on MC-LAG. The operation of dual homing at BSA level is based on two mechanisms. Ring control connection between two BSAs have two components, in-band and out-of-band communication. With in-band communication, BFD session between BSA-1 and BSA-2 running through the access ring and using dedicated IES/VPRN interface configured on both nodes. This connection uses a separate VLAN throughout the ring. The access nodes provides transparent bridging for this VLAN. The BFD session is used to continuously verify the integrity of the ring and to detect a failure somewhere in the ring.

With out-of-band communication, the communication channel is used by BSA nodes to exchange information about the reachability of individual access nodes as well as basic configurations in order to verify the consistency of the ring. The configuration information is synchronized through multi-chassis synchronization and therefore it is mandatory to enable multi-chassis synchronization between two nodes using the multi-chassis-ring concept.

In addition, the communication channel used by MC-LAG or MC-APS control protocol is used to exchange some event information. The use of this channel is transparent to the user.

Ring node connectivity check continuously checks the reachability of individual access nodes in the ring. The session carrying the connection is conducted on separate VLAN, typically common for all access nodes. SHCV causes no interoperability problems.

9.3.20.2 Steady-State Operation of Dual-homed Ring

Figure 122 illustrates the operation of the dual-homed ring. The steady state is achieved when both nodes are configured in a consistent way and the peering relation is up. The multi-chassis ring must be provisioned consistently between two nodes.

In-Band Ring Control Connection (IB-RCC) is in operational UP state. Note that this connection is set up using a bi-directional forwarding session between IP interfaces on BSA-1 and BSA-2.

Figure 122 **Dual Homing Ring Under Steady-State Condition**

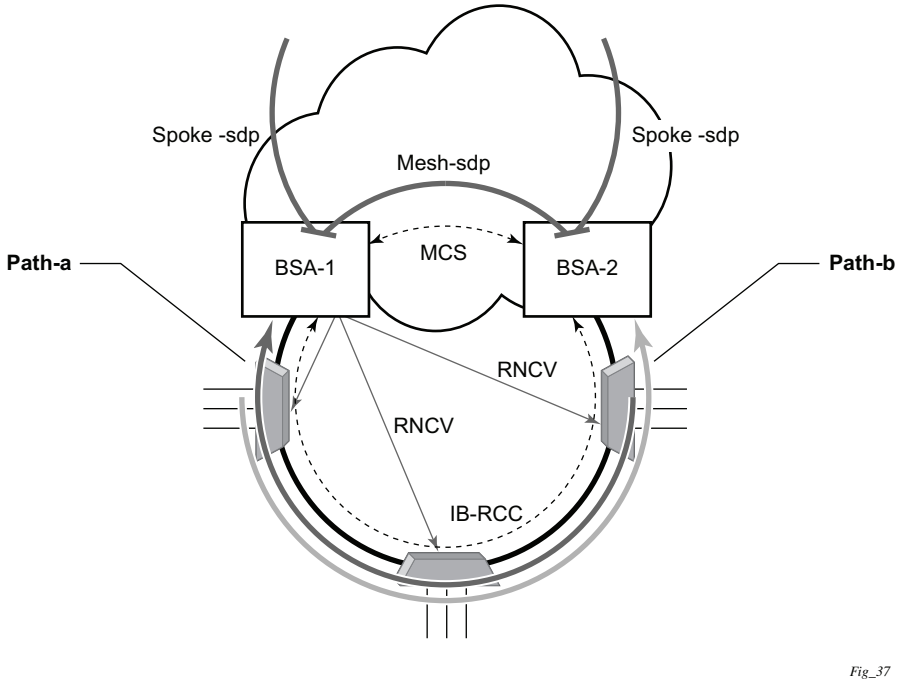

In Figure 122, the ring is fully closed and every access node has two possible paths towards the VPLS core. Figure 122 refers to these as **path-a** and **path-b**. In order to avoid the loop created by the ring, only one of the paths can be used by any given ring node for any given VLAN. The assignment of the individual VLANs to path-a or path-b, respectively, has to be provisioned on both BSAs.

The selection of the BSA master for both paths will be based on the IP address of the interface used for IB-RCC communication (bi-directional forwarding session). The BSA with the lower IP address of the interface used as IB-RCC channel will become master for ring nodes and their respective VLANs assigned to path-a. The master of path-b will be other BSA.

In this example, each path in the ring has a master and standby BSA. The functionality of both devices in steady state are as follows:

In the master BSA:

- All SAPs that belong to the path where the given BSA is a master, are operationally UP and all FDB entries of subscriber hosts associated with these SAPs point to their respective SAPs.
- The master of a path performs periodical Ring Node Connectivity Verification (RNCV) check to all ring nodes.

- In case of a RNCV failure, the respective alarm will be raised. Note that the loss of RNCV to the given ring node does not trigger any switchover action even if the other BSA appears to have the connection to that ring node. As long as the BFD session is up, the ring is considered closed and the master/standby behavior is driven solely by provisioning of the individual paths.
- The ARP reply agent replies to ARP requests addressing subscriber hosts where the BSA master.

In the standby BSA:

All SAPs that belong to a BSA's path, the standby will be operationally down and all FDB entries of subscriber hosts associated with those SAPs will be pointing towards SDP connecting to master BSA (also called a shunt SDP).

In both BSAs:

- The information on individual paths assignment is exchanged between both BSAs through multi-chassis synchronization communication channel and conflicting SAPs (being assigned to different paths on both BSA nodes) will be forced to path-a (the default behavior).
- For IGMP snooping, the corresponding multi-chassis IDs are targeting all subscriber-facing SAPs on both nodes. On the standby BSA node, the corresponding SAPs are in an operationally down state to prevent the MC traffic be injected on the ring twice.

9.3.20.3 Broken-Ring Operation and the Transition to this State

Figure 123 illustrates the model with a broken ring (link failure or ring node failure). This state is reached in following conditions:

- Both nodes are configured similarly.
- Peering is up.
- The multi-chassis ring is provisioned similarly between two nodes
- IB-RCC is operationally down.

In this scenario, every ring node has only one access path towards the VPLS core and hence, the Path-a and Path-b notion has no meaning in this situation.

Functionally, both BSAs are now the master for the reachable ring nodes and will take action as described in Steady-State Operation of Dual-homed Ring. For all hosts behind the unreachable ring nodes, the corresponding subscriber host FDB entries point to the shunt SDP.

Figure 123 **Broken Ring State**

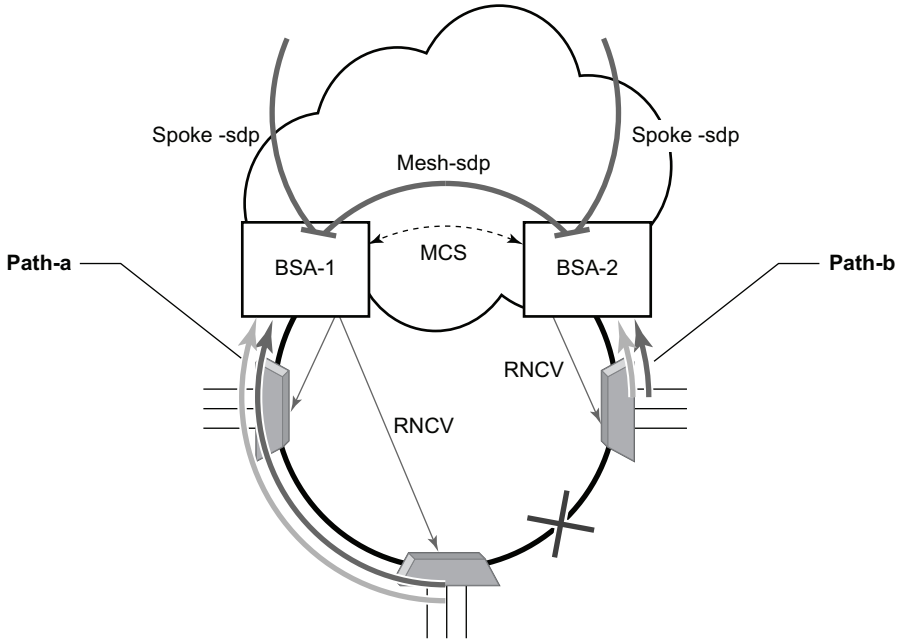

The mapping of individual subscriber hosts into the individual ring nodes is complicated, especially in the VLAN-per-service model where a single SAP can represent all nodes on the ring. In this case, a given BSA can have subscriber hosts associated with the given SAP that are behind reachable ring nodes as well as subscriber hosts behind un-reachable ring nodes. This means that the given SAP cannot be placed in an operationally down state (as in a closed ring state), but rather, selectively re-direct unreachable subscriber states to the shunt SDP.

All SAPs remain in an operationally up state as long as the ring remains broken. This mainly applies for BTV SAPs that do not have any subscriber hosts associated with and do not belong to any particular ring node.

In order to make the mapping of the subscriber-hosts on the given ring node automatically provisioned, the ring node identity will be extracted during subscriber authentication process from RADIUS or from a Python script. The subscriber hosts which are mapped to non-existing ring node will remain attached to the SAP.

At the time both BSA detect the break in IB-RCC communication (if BFD session goes down) following actions are taken:

- Both nodes trigger a RNCV check towards all ring nodes. The node, which receives the reply first, will assume a master functionality and will inform the other BSA through an out-of-band channel. This way, the other node can immediately take actions related to the standby functionality without waiting for an RNCV timeout. Even if the other node receives an RNCV response from the given ring node later, the master functionality remains with the node the received the response first.
- After assuming the master functionality for hosts associated with the given SAP(s), the node will send out FDB population messages to ensure that new path towards the VPLS core is established. The FDB population messages are sourced from the MAC address of the default gateway used by all subscriber hosts (such as the VRRP MAC address) which is provisioned at the service level.

9.3.20.4 Transition from Broken to Closed Ring State

By its definition, the multi-chassis ring operates in a revertive mode. This means that whenever the ring connectivity is restored, the BSA with lower IP address in the IB-RCC communication channel will become master of the path-a and vice versa for path-b.

After restoration of BFD session, the master functionality, as described in Steady-State Operation of Dual-homed Ring, is assumed by respective BSAs. The FDB tables are updated according to the master/standby role of the given BSA and FDB population messages is sent accordingly.

9.3.20.5 Provisioning Aspects and Error Cases

The multi-chassis ring can operate only if both nodes similarly configured. The peering relation must be configured and both nodes must be reachable at IP level. The multi-chassis ring with a corresponding sync-tag as a ring-name identifying a local port ID must be provision on both nodes. And, BFD session and corresponding interfaces needs to be configured in a consistent way.

In case the multi-chassis rings are not provisioned consistently, the ring will not become operational and the SAP managed by it will be in operationally up state on both nodes.

The assignment of individual SAPs to path-a and path-b is controlled by configuration of VLAN ranges according to the following rules:

- By default, all SAPs (and hence all VLANs on the given port) are assigned to path-a.
- An explicit statement defining the given VLAN range assigns all SAPs falling into this range to the path-b.
- An explicit statement defining the given VLAN range defines all SAPs that are excluded from the multi-chassis ring control.
- In case of a conflict in the configuration of VLAN ranges between two redundant nodes is detected, all SAPs falling into the “conflict-range” will be assigned to path-a, on both nodes regardless the local configuration.
- For QinQ-encapsulated ports the VLAN range refers to the outer VLAN.

9.3.20.6 Dual Homing to Two BSR Nodes

Figure 124 depicts a single DSLAM dual-homed to two BSRs.

Figure 124 **Low**

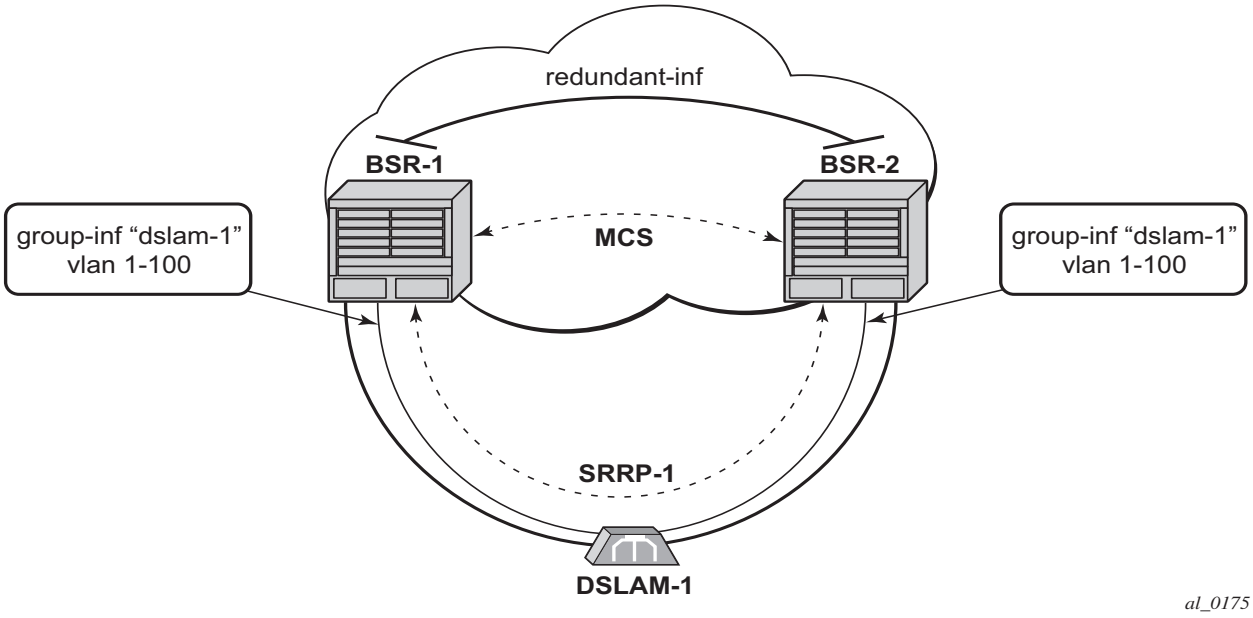

In order to provide dual-homing in the context of subscriber interfaces, the following items must be configured on both BSRs:

- Group interface (dslam-1) with corresponding SAPs (vlan 1-100)
- SRRP instance controlling given group interface
- Redundant interface between BSRs to provide “shunt” connectivity

- MCS connection to provide synchronization of dynamic subscriber-host entries

During the operation, BSR-1 and BSR-2 resolves master-backup relations and populates respective FDBs in such a way that at master side, subscriber-host entries point to corresponding group-interface while at the back-up side, subscriber-host entries point to the redundant interface. Note that the logical operation of the ring in the Layer 3 CO model is driven by SRRP. For more details on SRRP operation, see the Subscriber Routed Redundancy Protocol (SRRP) chapter.

9.3.20.7 MC Services

The typical implementation of MC services at the network level is shown in Figure 125.

Figure 125 **MC Services in a Layer 3-Ring Topology (a)**

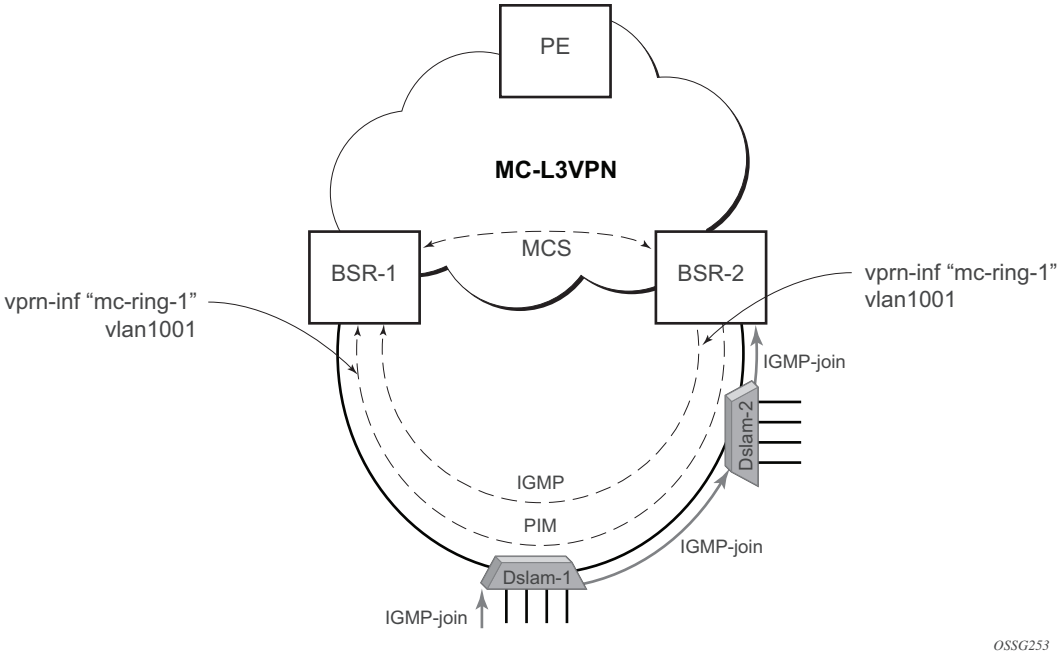

The IGMP is used to register joins and leaves of the user. IGMP messaging between BSRs is used to determine which router performs the querier role (BSR2 in Figure 125). PIM is used to determine which router will be the designated router and the router that sends MC streams on the ring.

The access nodes have IGMP snooping enabled and from IGMP messaging between BSR, they are aware which router is the querier. In the most generic case, IGMP snooping agents (in access nodes) send the IGMP-joins messages only to IGMP-querier. The synchronization of the IGMP entries can be then be performed through MCS. In some cases, access nodes can be configured in such a way that both ring ports are considered as m-router ports and IGMP joins are sent in both directions.

All of the above is a steady state operation which is transparent to the topology used in a Layer 2 domain.

A ring-broken state is shown in Figure 126.

In this case, IGMP and PIM messaging between BSRs is broken and both router assume role of querier and role of designated router. By the virtue of ring topology, both routers will see only IGMP joins/leaves generated by host attached to a particular “half” of the ring. This means that both routers will have “different” view on dynamic IGMP state.

Figure 126 **MC Services on a Layer 3-Ring Topology (b)**

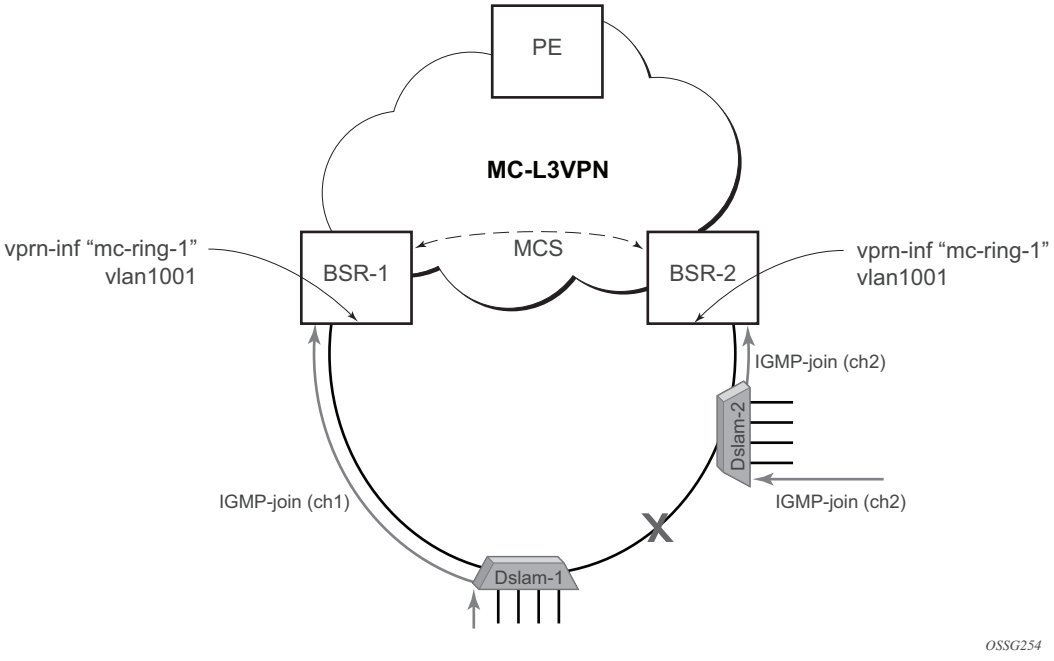

In principle, MCS could be used to synchronize both routers, but in case of a Layer 2 ring, the implementation sends all IGMP messages to a “ring-master: which then performs IGMP processing and consequently, MCS sync. As a result, any race conditions are avoided.

Another ring-specific aspect is related to ring healing. The ring continuity check is driven by BFD which then drives SRRP and PIM messaging. BFD is optimized for fast detection of ring-down events while ring-up events are announced more slowly. There is a time window when routers are not aware that the ring is recovered. In the case of MC, this means traffic will be duplicated on the ring.

To avoid this, the implementation of BFD provides a “raw mode” which provides visibility on “ring-up” events. The protocols, such as SRRP and PIM, use this raw mode rather than the BFD API.

9.3.20.8 Routed CO Dual Homing

Routed CO dual homing is a solution that allows seamless failover between nodes for all models of routed CO. In the dual homed environment, only one node will forward downstream traffic to a given subscriber at a time. Dual homing involves several components:

- Redundant Interface — This is used to shunt traffic to the active node for a given subscriber for downstream traffic.
- SRRP — This is used to monitor the state of connectivity to the DSLAM. See the SRRP section for more detail.
- MCS — This is used to exchange subscriber host and SRRP information between the dual homed nodes.

Routed CO dual homing can be configured for both wholesaling models. Dual homing is configured by creating a redundant interface that is associated with the protected group interfaces. The failure detection mechanism can be VRRP. If VRRP is used, each node monitors the VRRP state to determine the priority of its own interface.

Dual homing is used to aggregate a large number of subscribers in order to support a redundancy mechanism that will allow a seamless failover between nodes. Because of the Layer 3 nature of the model, forwarding is performed for the full subscriber subnet.

9.3.20.8.1 Redundant Interfaces

In dual homing, a redundant interface must be created. A redundant interface is a Layer 3 spoke SDP-based interface that allows delivery of packets between the two nodes. The redundant interface is required to allow a node with a failed link to deliver packets destined to subscribers behind that link to the redundant node. Since subscriber subnets can span multiple ports it is not possible to stop advertising the subnet, thus, without this interface the node would black hole.

The redundant interface is associated with one or more group interfaces. An interface in backup state will use the redundant interface to send traffic to the active interface (in the active node). The SAP structure under the group interface must be the same on both nodes as the synchronization of subscriber information is enabled on a group interface basis. Traffic can be forwarded through the redundant interface during normal operation even when there are no failed paths. See Figure 127.

9.3.20.8.2 SRRP in Dual Homing

Subscriber Router Redundancy Protocol (SRRP) allows two separate connections to a DSLAM to operate in an active/standby fashion similar to how VRRP interfaces operate. Since the SRRP state is associated with the group-interface, multiple group-interfaces may be created for a given port such that some of the SAPs will be active in one node and others active on the other node. While each SRRP pair is still allowed to be active/backup, the described configuration is allowed for load balancing between the nodes. Note that in a failure scenario subscriber bandwidth will be affected. For more information about SRRP, see the Subscriber Routed Redundancy Protocol (SRRP) chapter.

If SRRP is configured before the redundant interface is up, and in backup state the router will forward packets to the access node via the backup interface but will not use the gateway MAC address. This applies to failures in the redundant interface as well. If the redundant interface exists and up the router will send downstream packets to the redundant interface and will not use the backup group interface.

In a dual homing architecture the nodes must be configured with SRRP to support redundant paths to the access node. The nodes must also be configured to synchronize subscriber data and IGMP state. To facilitate data forwarding between the nodes in case some of the ports in a given subscriber subnet are affected a redundant interface must be created and configured with a spoke. The redundant interface is associated with one or more group interfaces.

The service IDs for both the wholesale VPRN and the retailer VPRN must be the same in both nodes.

An interface in backup state will use the redundant interface to send traffic to the active interface (in the active node). The SAP structure under the group interface must be the same on both nodes as the synchronization of subscriber information is enabled on a group interface basis.

SRRP is associated a group-interface. Multiple group-interfaces can be created for a given port such that some of the SAPs will be active in one node and others active on the other node. While every SRRP pair is still allowed to be active/backup the described configuration will allow for load balancing between the nodes. Note that in a failure scenario subscriber bandwidth will be affected.

Figure 127 **Dual Homing Example**

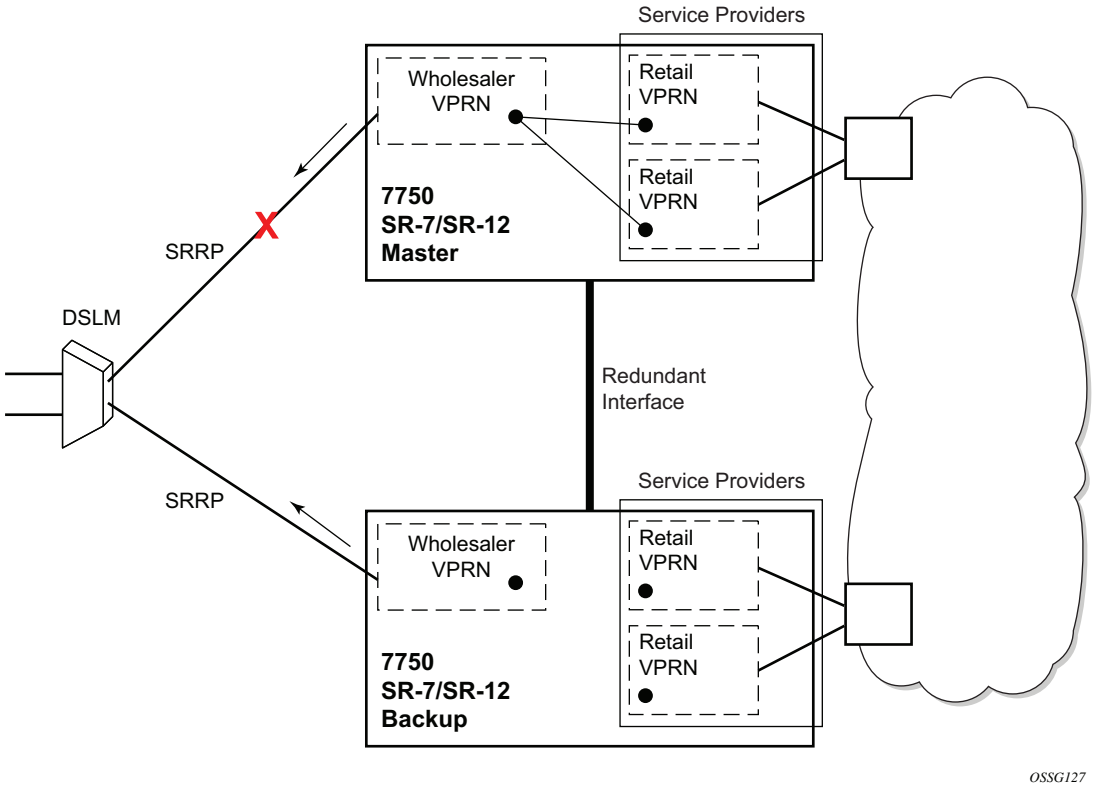

9.3.20.8.3 Synchronization

To establish subscriber state the nodes must synchronize subscriber information. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Basic System Configuration Guide* for multi-chassis synchronization configuration information. The operator must complete the configuration and the system must have data synchronized before the backup node may deliver downstream packets to the subscriber.

If dual homing is used with regular interfaces that run IGMP the nodes must be configured to synchronize the Layer 3 IGMP state.

The service IDs for both the wholesale VPRN and the retailer VPRN must be the same in both nodes.

9.3.20.8.4 Wholesale-Retail Multi-Chassis Redundancy

Multi-Chassis Redundancy for a retail service is enabled with the SRRP and redundant interface configuration on the wholesale group-interface parented by the forwarding subscriber interface. The Multi-Chassis state (active or standby) of the retail subscriber host is determined from the SRRP state (master / non-master) of the group-interface that parents the SAP of the retail subscriber host. The retail service id must be equal on both nodes.

Sample wholesale service configuration:

```
vprn 3000 customer 1 create
           description "Wholesale service"
           route-distinguisher 64500:3000
           auto-bind-tunnel
               resolution-filter
                   ldp
                   rsvp
               exit
               resolution filter
           exit
           vrf-target import target:64500:3000
           redundant-interface "red-int-1" create
               address 192.168.100.0/31
               ip-mtu 1500
               spoke-sdp 12:3000 create
                   no shutdown
           exit
           subscriber-interface "sub-int-1" create
               address 10.1.1.253/24 gw-ip-address 10.1.1.254
               address 10.1.2.253/24 gw-ip-address 10.1.2.254
               group-interface "group-int-1-1" create
                   dhcp
                       --- snip ---
                   exit
                   redundant-interface "red-int-1"
                   sap 1/1/6:1.4001 create
                       description "SRRP 1 message path"
                   exit
                   srrp 1 create
                       message-path 1/1/6:1.4001
                       send-fib-population-packets outer-tag-only
                       no shutdown
                   exit
                   pppoe
                       --- snip ---
                   exit
```

```
                exit
                group-interface "group-int-1-2" create
                    dhcp
                        --- snip ---
                    exit
                    redundant-interface "red-int-1"
                    sap 1/1/6:2.4001 create
                        description "SRRP 2 message path"
                    exit
                    srrp 2 create
                        message-path 1/1/6:2.4001
                        priority 50
                        send-fib-population-packets outer-tag-only
                        no shutdown
                    exit
                    pppoe
                        --- snip ---
                    exit
                exit
            exit
            no shutdown
        exit
```

Sample retail service configuration:

```
 vprn 3001 customer 1 create
            description "Retail service"
            route-distinguisher 64500:3001
            auto-bind-tunnel
                resolution-filter
                    ldp
                    rsvp
                exit
                resolution filter
            exit
            vrf-target target:64500:3001
            subscriber-interface "sub-int-rt-3000-1" fwd-service 3000 fwd-
subscriber-
interface "sub-int-1" create
                address 10.1.11.253/24 gw-ip-address 10.1.11.254
                address 10.1.12.253/24 gw-ip-address 10.1.12.254
                dhcp
                    --- snip ---
                exit
                pppoe
                    --- snip ---
                exit
            exit
            no shutdown
        exit
```

Overlapping addresses in retail subscriber interfaces (enabled with **private-retail-subnets**) cannot be used in combination with multi-chassis redundancy. This is because retail subscriber subnets and prefixes leaked in the wholesale service are needed to forward downstream shunted traffic over the redundant interface. For the same reason, retail unnumbered host routes must be leaked in the wholesale service:

- The address of an IPv4 unnumbered subscriber host (enabled with **unnumbered** *ip-int-name* | *ip-address* or **allow-unmatching-subnets** on the retail subscriber interface) is not contained in the subnets configured on the retail subscriber interface. The export of the IPv4 retail subscriber host routes to the wholesale service must be explicitly enabled with the **export-host-routes** command:

```
        vprn 3001 customer 1 create
            subscriber-interface "sub-int-rt-3000-1" fwd-service 3000 fwd-
subscriber-
interface "sub-int-1" create
                allow-unmatching-subnets
                address 10.1.11.253/24 gw-ip-address 10.1.11.254
                address 10.1.12.253/24 gw-ip-address 10.1.12.254
                export-host-routes
                --- snip ---
```

- The address of an IPv6 unnumbered subscriber host (enabled with **ipv6 allow-unmatching-prefixes** on the retail subscriber interface) is not contained in the IPv6 prefixes configured on the retail subscriber interface. IPv6 retail subscriber host routes are automatically exported to the wholesale service.

Downstream traffic arriving on an SRRP non-master node must be shunted over the redundant interface. Downstream traffic shunting can be reduced by advertising the retail subscriber subnets and prefixes from the master SRRP node with a more favorable metric using routing policies. To make retail subscriber subnets and prefixes SRRP state-aware, they have to be configured to track an SRRP instance that is active on a group interface of the connected wholesale subscriber interface:

```
        vprn 3001 customer 1 create
            subscriber-interface "sub-int-rt-3000-1" fwd-service 3000 fwd-
subscriber-
interface "sub-int-1" create
                address 10.1.11.253/24 gw-ip-address 10.1.11.254 track-srrp 1
                address 10.1.12.253/24 gw-ip-address 10.1.12.254 track-srrp 2
                ---snip---
                ipv6
                    subscriber-prefixes
                        prefix 2001:db8:d001::/48 pd track-srrp 1
                        prefix 2001:db8:d002::/48 pd track-srrp 2
                        prefix 2001:db8:1:100::/56 wan-host track-srrp 1
                        prefix 2001:db8:1:200::/56 wan-host track-srrp 2
                    exit
                exit
```

Multi-Chassis Redundancy is supported for IPoE (IPv4 and IPv6) and PPPoE (IPv4 and IPv6) retail subscriber hosts and sessions.

9.3.20.9 SRRP and Multi-Chassis Synchronization

In order to take full advantage of SRRP resiliency and diagnostic capabilities, the SRRP instance will be tied to a MCS peering that terminates on the redundant node. Once the peering is associated with the SRRP instance, MCS will synchronize the local information about the SRRP instance with the neighbor router. MCS automatically derives the MCS key for the SRRP instance based on the SRRP instance ID. An SRRP instance ID of 1 would appear in the MCS peering database with a MCS-key srrp-0000000001.

The SRRP instance information stored and sent to the neighbor router consists of:

- The SRRP instance MCS key
- Containing service type and ID
- Containing subscriber IP interface name
- Subscriber subnet information
- Containing group IP interface information
- The SRRP group IP interface redundant IP interface name, IP address and mask
- The SRRP advertisement message SAP
- The local system IP address (SRRP advertisement message source IP address)
- The group IP interface MAC address
- The SRRP gateway MAC address
- The SRRP instance administration state (up/down)
- The SRRP instance operational state (disabled/becoming-backup/backup/ becoming-master/ master)
- The current SRRP priority
- Remote redundant IP interface availability (available/unavailable)
- Local receive SRRP advertisement SAP availability (available/unavailable)

9.3.20.10 Dual Homing and ANCP

The routers provide a feature related to exchange of control information between DSLAM and BRAS (BSA is described in this model). This exchange of information is implemented by in-band control connection between DSLAM and BSA, also referred to as ANCP connection.

In case of dual homing, two separate connections will be set. As a consequence, there is no need to provide synchronization of ANCP state. Instead every node of the redundant-pair obtains this information from the DSLAM and creates corresponding an ANCP state independently.

9.3.21 SRRP Enhancement

The SRRP enhancements addressed in this section is to reduce the need for redundant-interface between the pair of redundant nodes without sacrificing the subnet aggregation on the back-end.

Redundant BNG nodes are not always collocated. This means that the logical link associated with the redundant (shunt) interfaces is taking the uplink path thus wasting valuable bandwidth (downstream traffic that arrives to the Standby node is routed via uplinks for the second time over to the Master node).

To meet the requirement to reduce the existence of shunted traffic only to the short transitioning period between SRRP switchovers while the routing on the network side is converging, the following was required (referring to Figure 128):

1. Share IP subnets over multiple SRRP instances. This is not mandatory, but it would help to load balance traffic over the two nodes. For example, IP subnets 10 and 11 can be shared over SRRP instances 10 and 20 on node 1, and the IP subnet 12 can be associated with the SRRP instance 30 on node 2.
2. *SRRP aware routing* – this allows to dynamically increase routing metric on the IP subnets advertised from the Master SRRP node in comparison to the Standby SRRP node. It also allows to advertise/withdraw routes from a routing protocol based on the SRRP state. In this fashion, downstream traffic is routed in a predictable manner towards the Master SRRP node.
3. *SRRP Fate Sharing* for SRRP instances 10 and 11. This ensures congruency of SRRP states on the same node. This is a necessary step towards *SRRP aware routing*.

Figure 128 **IP Subnet Per SRRP Master Group**

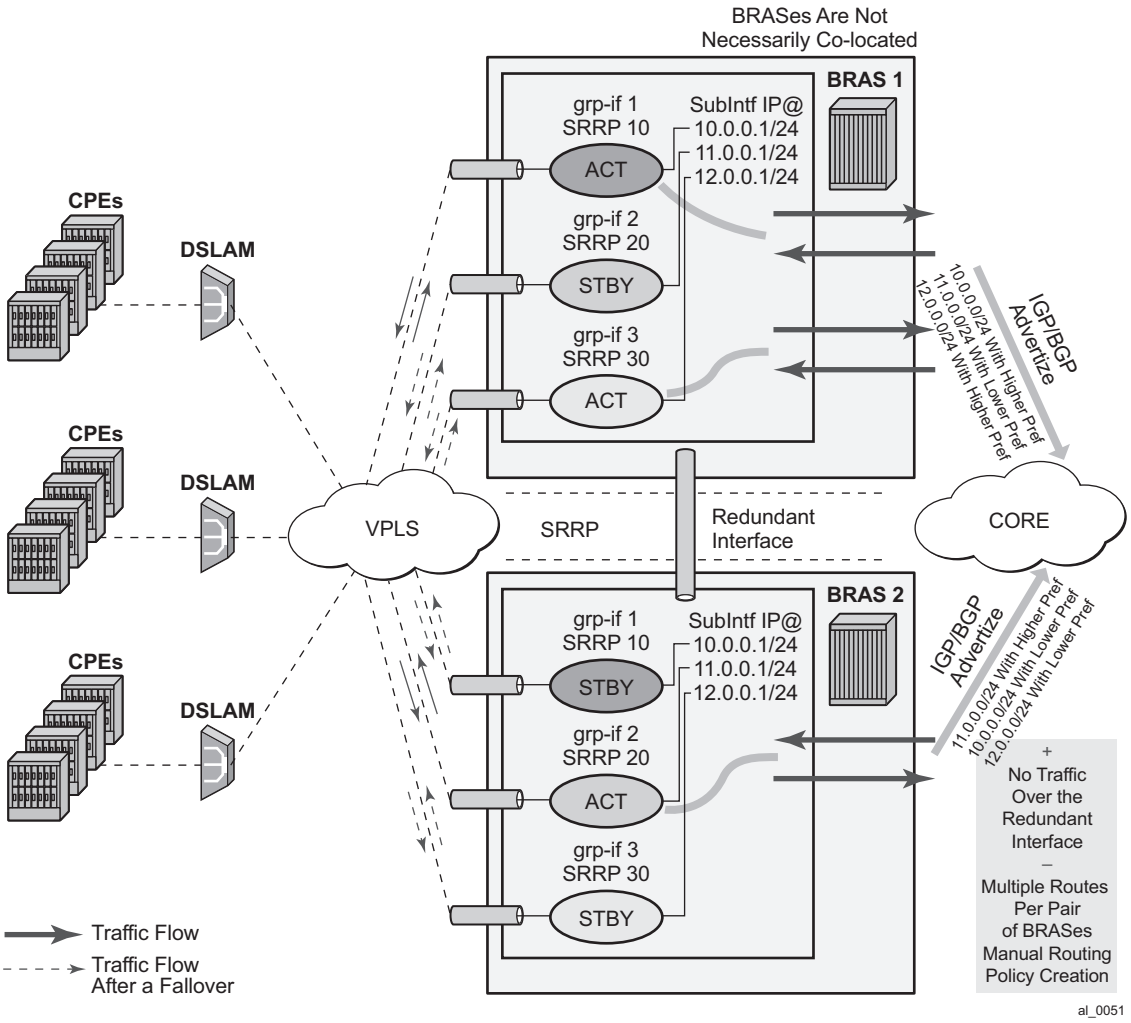

9.3.21.1 SRRP Fate Sharing

SRRP Fate Sharing is a concept in which a group of SRRP instances track a single operational-object comprised of SRRP messaging SAPs. The SRRP instances behave as one (in the single failure case) with regards to SRRP mastership. The group of SRRP instances that are sharing fate on a paired node are referred as a Fate Sharing Group (FSG).

Transition of a single messaging SAP within the FSG into a DOWN state forces the SRRP instance on top of it into the INIT state. Consequently, all other SRRP instances within the same FSG transitions into a Backup state. In other words, SRRP instances within the FSG all share the same fate as the failed SRRP instance as shown in Figure 129. SRRP Fate Sharing provides optimal protection in the context of a single failure in the network.

Figure 129 **FSG — Single Network Failure**

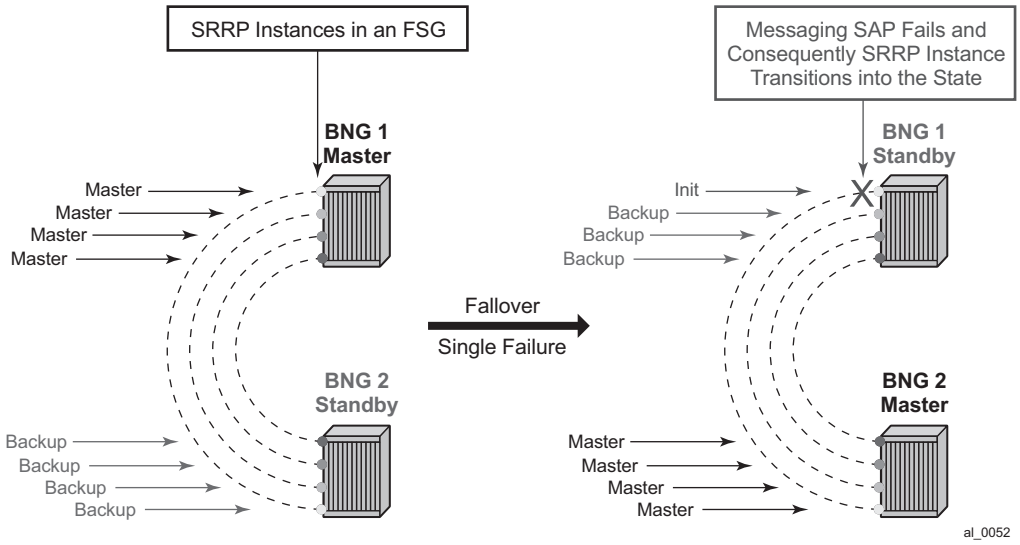

In the case of multiple network failures, the concept of the FSG breaks as there is a possibility that a ‘FSG’ contains SRRP instances that are in any of the three possible SRRP states: Master, Backup, or Init. This Fate Sharing feature may not provide optimal protection when there are multiple network failures distributed over both redundant nodes.

Figure 130 **Multiple Network Failures**

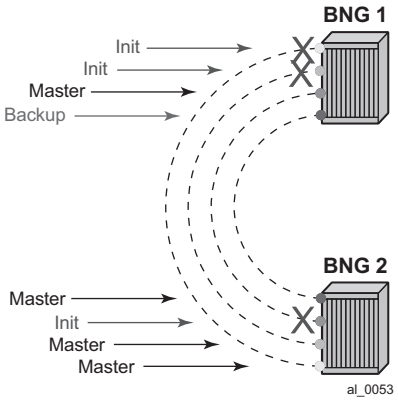

The whereabouts of the failure in the network path that SRRP is designed to monitor are not always clearly reflected through SRRP states. For example, if the network failure is somewhere in the aggregation network beyond the direct reach of our BNG, SRRP assumes Mastership on both BNG nodes. This is a faulty condition and the reason why solely monitoring of the SRRP states is not enough to protect against failures. On the other hand, the SRRP messaging SAP states are more indicative of the network failure since they can be tied into Eth-OAM.

Once a single network failure is detected and as a result an SRRP instance transitions into a non-Master state, the remaining SRRP instances in the FSG are forced into a Backup state. This is achieved by changing the priority of each individual SRRP instance in the FSG.

In the case of simultaneous multiple failures (multiple ports fail at the same time), it is possible that the SRRP instances within the FSG settle in any of the three possible SRRP states: Master, Backup, or Init. In such scenario shunted traffic will ensue.

In the premise of SRRP Fate Sharing, the network failure will be reflected in the operational state of the messaging SAP over which SRRP runs. This will certainly be the case if the failure is localized to the BNG (somewhere on the directly connected link). In the case of non-localized failure (beyond the direct reach of the BNG node), Eth-OAM might be needed in to detect the remote end failure and consequently bring the SAP operationally into a DOWN state.

Once the single network failure is detected, all instance within the FSG transitions into a non-Master state.

If there are no failures in the network, all SAPs are UP and SRRP instances within the FSG are in a homogeneous and deterministic state based on their configured priorities.

Figure 131 **SRRP Fate Sharing**

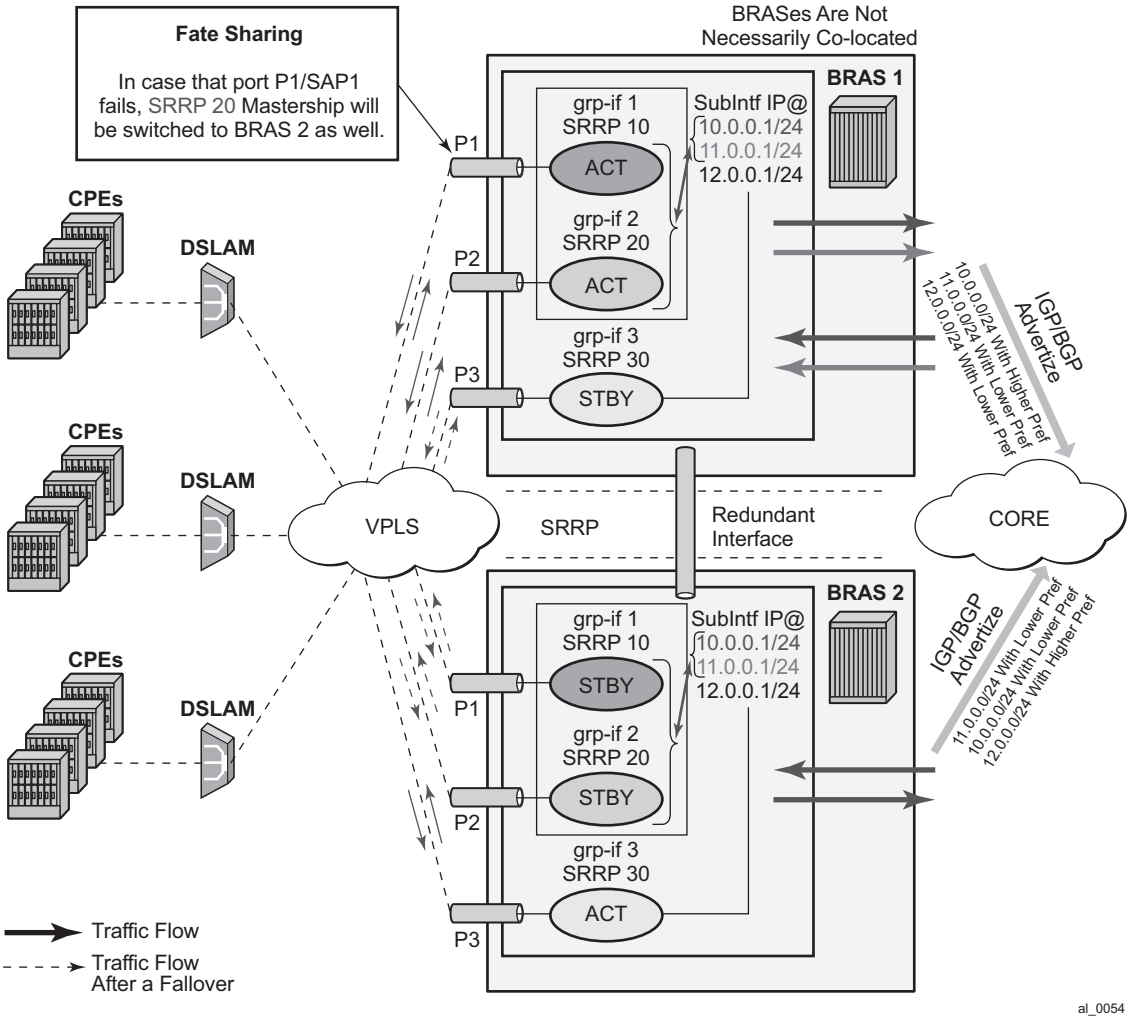

Failure Detection in a Fate Sharing Group

1. Dual homing over directly connected ports.
   No Eth-OAM is needed, AN is directly connected to the BNG.

Figure 132 **Scenario 1**

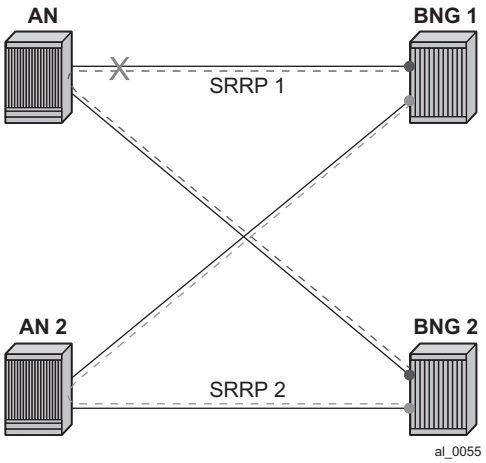

2. Dual homing with aggregation network - aggregation network has no redundancy between Layer 2 switches (STP). To determine whereabouts of failure at point 1 in the figure below, Eth-OAM is needed.

Figure 133 **Scenario 2**

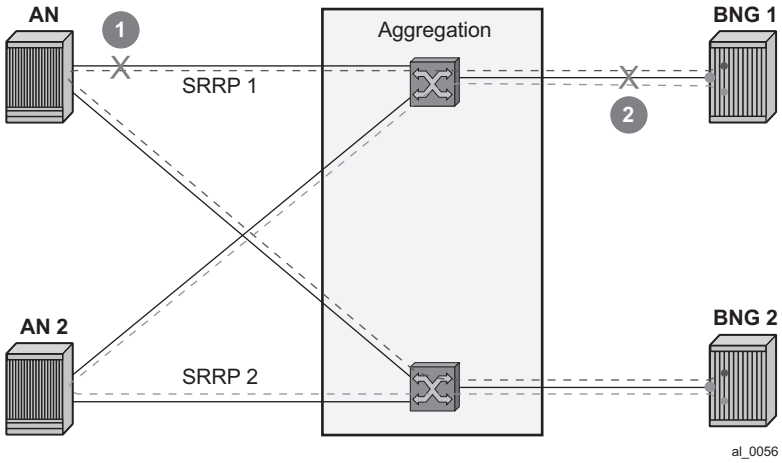

3. Dual homing with aggregation network - aggregation network with redundancy between Layer 2 switches (STP).
   No Eth-OAM is needed in this case for successful operation. However, the failure detection is based on the failure of the directly attached ports.

Figure 134 **Scenario 3**

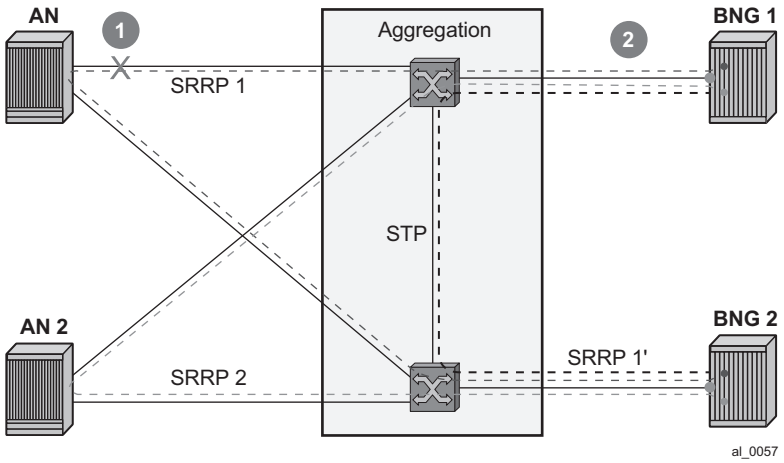

4. Single homing with aggregation network.
   In this case, SRRP can protect only against direct failures. Any remote failure leaves a part of the network isolated from the subscriber point of view.

Figure 135 **Scenario 4**

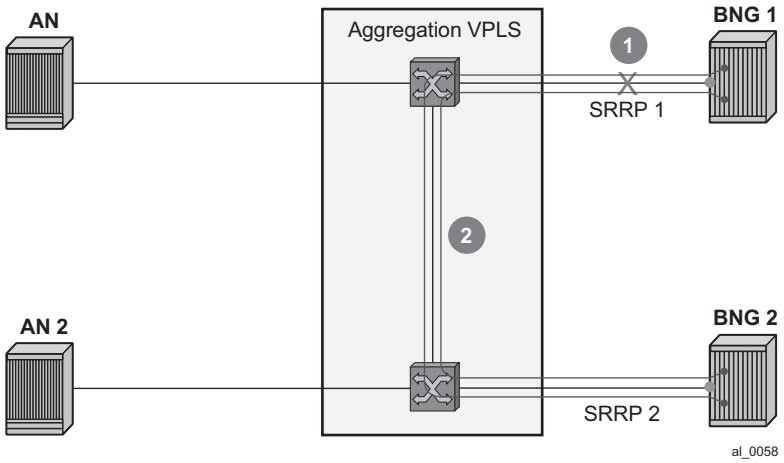

9.3.21.2 Fate Sharing Algorithm

Fate Sharing Algorithm (FSG) is relaying on tracking the state of messaging SAPs over which SRRP instances run. An SRRP instance with the messaging SAP operationally DOWN will transition into the Init state.

The transitioning of any messaging SAP in a FSG into an UP/DOWN state will trigger SRRP priority adjustment within the FSG. The SRRP priorities should be chosen carefully to achieve the desired behavior. They are modified dynamically as the SAP states change. The range in which SRRP priorities can be modified is from 1 to the SRRP priority that is initially configured under the SRRP node. Here are some general guidelines for choosing SRRP priorities in a FSG:

- Initially configured SRRP priorities for all SRRP instance within the FSG within the node should be the same.
- Initially configured SRRP priorities should be different between pairing FSGs. For example, SRRP instances in the BNG node A within an FSG will all have the same SRRP priority ‘X’, while corresponding SRRP instances on the paired node within corresponding FSG will all have SRRP priority ‘Y’. This ensures that SRRP mastership is clearly defined between the two BNG nodes. Note that this step is not mandatory as SRRP will naturally break the Mastership tie in the case that all SRRP priorities are the same. However, following this step may provide a clearer view from an operational perspective.
- The priority-step used for dynamic SRRP priority adjustment must be greater than the difference in initially configured SRRP priorities between two BNG nodes. This ensures that a single failure event triggers the SRRP switchover. Otherwise, if the dynamically lowered SRRP priority is still greater than the one from the SRRP peer, the switchover would not be triggered. Therefore, the fate sharing concept would not function as intended.
- Initially configured SRRP priority of each SRRP instance should be greater than the (anticipated) number of SRRP instances in a FSG multiplied by the SRRP priority-step. This ensures that the dynamically priority never tries to go below 1. There is a code check that prevents SRRP priority going below 1. Nonetheless, it is recommended not to get into a situation where this needs to be enforced in the code.

The priorities will never be less than 1 or greater than initially configured SRRP priority.

Example scenarios:

Assume 3 SRRP instances in a FSG. The SRRP instances in the FSG-1 on BNG 1 have the priority of 100, while the SRRP instances in the FSG-2 on BNG 2 have the priority of 95. The priority-step is 6. The SRRP instances and underlying messaging SAPs will be referred as SRRP 1, 2, 3 and SAP 1,2,3, respectively.

Initialization:

Scenario 1 – all SAPs are operationally UP.

BNG 1 boots up and all messaging SAPs transition into the UP state. When the first SRRP instance in FSG-1 comes up, it looks under the FSG to finds out how many messaging SAPs are operationally UP. Since all messaging SAPs are operationally UP, this first SRRP instance assumes its initially configured priority of 100. The other two SRRP instances in the same FSG follows the same sequence of events.

BNG 2 follows the same flow of events. As a result, BNG 1 assumes mastership over BNG 1 for all SRRP instances within the corresponding FSG.

Scenario 2 – messaging SAP 1 is operationally DOWN on BNG 1, the rest of the messaging SAPs are operationally UP.

SRRP 2 and 3, during the initialization, pick up SRRP priority of 94 (100 – 1*priority-step).

On BNG 2, all messaging SAPs are UP and consequently all SRRP instances within the FSG on BNG 2 have SRRP priority of 95. The SRRP instances on BNG 2 assumes Mastership.

Scenario 3 – Continuing from scenario 2, the SAP 1 on BNG 1 transitions into the UP state. SRRP priority of each SRRP instance in FSG-1 is increased by 6, bringing it to 100, enough to assume Mastership.

Adding a New Instance into an FSG

To introduce minimal network disruption, first create messaging SAPs in both BNG nodes and ensure that both SAPs are operationally UP. Then a new SRRP 4 instance should be created on both BNG nodes. The next step would be to include this new messaging SAP into a SAP monitoring group. And finally, the SRRP-4 is added into the FSG (1 and 2). This triggers the recalculation of SRRP priorities for the existing FSG-1 and FSG-2. Since all SRRP priorities are at the max (initially configured priority), nothing changes.

There are more disruptive ways of adding an SRRP instance into a FSG. One such example would be in the case where SRRP priorities are not at their maximum (initially configured) priority. If an SRRP instance is first added into an FSG that is in a Backup state, this would increase the FSG priority and potentially cause a switchover. If the SRRP instances is then added in a FSG on the peer BNG (previously Master), the priority of this FSG would be increased again and the switchover would unnecessarily occur for the second time. The new SRRP instances, once operational, should always be added in the Master FSG first.

SRRP priority re-calculation within the FSG is triggered by the following events:

- SRRP initialization
- addition of a SAP under the monitoring group

- messaging SAP failure

This priority calculation looks into how many SAPs are in the DOWN state within the monitored SAP group. Based on this number, the priority is calculated as follows:

SRRP priority = configured-priority – priority-step * num_down_SAPs.

9.3.21.3 SRRP Aware Routing - IPv4/IPv6 Route Advertisement Based on SRRP State

There are three cases that need to be covered, each case with its own specifics:

- Subscriber Interface Routes (IPv4/IPv6)
- Managed Routes
- Subscriber Management Routes (/32 IPv4 hosts routes and IPv6 PD wan-host routes)

Depending on the route type, the action is to either modify the route metric based on the SRRP state that the route is tracking, OR to advertise/withdraw the route based on the SRRP state that the route is tracking. The action is defined in the routing policy and it is based on the new attributes with which the routes are associated.

To achieve a better granularity of the routes that are advertised, an origin attribute is added to the subscriber management routes (/32 IPv4 routes and IPv6 PD wan-host) with three possible values:

aaa

IPv4

subscriber-management /32 host routes that are originated through RADIUS framed-ip-address VSA other than 255.255.255.254. The 255.255.255.254 returned by the RADIUS indicates that the BNG (NAS) should assign an IP address from its own pool.

IPv6

subscriber-management routes that are originated through framed-ipv6-prefix (SLAAC), delegated-ipv6-prefix (IA_PD) or alc-ipv6-address (IA_NA) RADIUS attributes. This is valid for IPoE and PPPoE type host.

dynamic

IPv4

subscriber-management /32 host routes that are originated via DHCP server (local or remote) and also RADIUS framed-ip-address=255.255.255.254 (RFC 2865).

IPv6

subscriber-management routes that are assigned via local DHCPv6 server pools whose name is obtained through Alc-Delegated-IPv6-Pool (PD pool) and Framed-IPv6-Pool (NA pool) RADIUS attributes. This is valid for IPoE and PPPoE type hosts.

In addition, for IPoEv6 only, the pool name can be also obtained via ipv6-delegated-prefix-pool (PD pool) and ipv6-wan-address-pool (NA pool) from LUDB.

static

IPv4

subscriber-management /32 host routes that are originated via LUDB. This also covers RADIUS fallback category (RADIUS falls back to system-defaults or to LUDB).

IPv6

subscriber-management routes obtained from LUDB via ipv6-address (IA_NA) or ipv6-prefix (IA_PD). This is supported only for IPoE.

Overall, the following new route attribute is added:

state: srrp-master, srrp-non-master

The existing origin attribute is expanded to contain the following values:

origin: aaa, dynamic, static

These two attribute types are applied in the following fashion:

The state attribute is applied to all three route types: *subscriber interface routes, managed routes* and *subscriber management routes*. Each route listens to the SRRP state.

If an attribute is defined in the routing policy as a match condition (from statement) but the route itself does not have this attribute, the route is evaluated into a non-match condition.

The origin attribute is always applied only to subscriber management routes. No additional statement is needed to explicitly apply this attribute as it may be the case for the state attribute.

Every time there is a change in the attribute associated with the route, the route is re-evaluated in the RTM by the routing policy and corresponding action is taken.

9.3.21.3.1 Subscriber Interface Routes (IPv4 and IPv6)

Optimized routing and elimination of downstream shunt traffic during normal operation can be achieved by **statically** favoring the routes on the network side that are advertised with an increased metric by Master SRRP nodes.

The downside of this static approach is that during the port/card failure and consequently a SRRP switchover, the node with the failed port/card will continue to advertise routes with the same high metric as long as the subscriber interface is in the ‘UP’ state (or a single SAP under it). That is, the network side will not be aware of the switchover. It will continue to forward traffic to the standby node, and as a result, heavy shunt traffic will ensue. To effectively deal with this, the network side must be aware of the routing change that occurred in the access layer.

When failure is detected, the metric for the route is changed automatically based on the following configuration:

```
configure
    service <type> <id>
        subscriber-interface <intf-name>
            address <ip-address> gw-ip-address <gw-address> track-srrp <srrp-
inst> holdup-time <msec>
            ipv6
            subscriber-prefixes
                prefix <ipv6-prefix> pd track-srrp <srrp-id> holdup-time <msec>
                prefix <ipv6-prefix> wan-host track-srrp <srrp-id> holdup-
time <msec>


    policy-options
        begin
           policy-statement <name>
            entry 1
                from
                    protocol direct
                    state ‘srrp-master’
                    exit
               action accept
                metric set 100
                exit
            exit
            entry 2
                from
                    protocol direct
                    state ‘srrp-non-master’
                    exit
                action accept
                metric subtract 10
                exit
```

```
exit
entry 3
    from
        protocol direct
        exit
    action accept
    exit
exit
```

This configuration ensures that the route metric is changed for the subscriber interface routes based on the SRRP state while the other, non-subscriber directly attached routes are unaffected by SRRP.

The **Route Advertisement based on SRRP State** requirement is applicable to BGP (IPv4, IPv4-IP-VPN) and IGP.

The routing policy also provides the flexibility to prevent route advertisement (*action reject*) instead of changing the route metric.

Although this feature is designed to minimize or eliminate the use of the redundant-interface, it is important to note that the redundant-interfaces would still be used in the case of transient conditions. An example of such condition would be:

1. Messaging SAP Fails
2. SRRP switches over
3. Stale routing in the core is still in the effect while the metric is being propagated (or the route is being advertised/withdrawn). During this time, traffic is flowing over the redundant interface.
4. Network convergence is complete
5. Traffic in the network core is redirected to the new Master SRRP node

9.3.21.3.2 Managed Routes

Only the state attribute is applicable to managed routes, and only to the ones that are synchronized (static and RADIUS obtained – framed-route and framed-ipv6-route). The managed routes obtained via BGP are not synchronized and this feature is not applicable to them.

Based on the SRRP state, the managed route can be either advertised with a modified metric or be withdrawn altogether.

For example:

Managed routes that are tracking SRRP state are only advertised from the Master node and denied from Backup node. All other managed routes that are not tracking SRRP state are advertised regardless of the SRRP state.

```
policy-options
    begin
        policy-statement <name>
            entry 1
                from
                    protocol managed
                    state 'srrp-master'
                    exit
                action accept
                exit
            exit
            entry 2
                from
                    protocol managed
                    state 'srrp-non-master'
                    exit
                action reject
                exit
            exit
            entry 3
                from
                    protocol managed
                    exit
                action accept
                exit
            exit
```

9.3.21.3.3 Subscriber Management Routes (/32 IPv4 Host Routes, IPv6 PD WAN-Host Routes)

Both attributes (state and origin) are applicable to the subscriber management routes.

For Example:

A Service Provider wants to advertise only subscriber-management routes with the origin dynamic and AAA from the Master node. Routes with the LUDB origin are not advertised. Standby node is not advertising any /32 subscriber management routes.

```
policy-options
    begin
        policy-statement <name>
            entry 1
                from
                    origin dynamic
                    state 'srrp-master'
                    exit
            action accept
```

```
            exit
        exit
    exit
           entry 2
                from
                    origin dynamic
                    state 'srrp-master'
                    exit
                action accept
                exit
        exit
```

Default action is reject.

9.3.21.3.4 Activating SRRP State Tracking

The SRRP state tracking by routes is turned on only when desired.

For subscriber-interface routes (IPv4 and IPv6), this is performed explicitly.

```
subscriber-interface <intf-name>
    address <ip-address> gw-ip-address <gw-address> track-srrp <inst-name> holdup-
time <msec>
        ipv6
        subscriber-prefixes
            prefix <ipv6-prefix> pd track-srrp <srrp-id> holdup-time <msec>
            prefix <ipv6-prefix> wan-host track-srrp <srrp-id> holdup-time <msec>
```

For managed and subscriber management routes, this is explicitly enabled under the group interface:

```
group-interface <name>
    srrp-enabled-routing holdup-time <msec>
```

9.3.21.4 SRRP in Conjunction with a PW in ESM Environment – Use Case

In certain cases, subscriber traffic is terminated on the BNG via an Epipe. In this case, the subscriber traffic can be offloaded onto a plain Ethernet port via a VSM module (a 'loop') so that it can be terminated in ESM. Epipes can be configured in A/S configuration and terminated on two BNG nodes in multihomed environment.

In such multi-homed environment with Epipes and ‘loops’, the ESM itself would be detached from the Epipe, which brings the subscriber traffic to the BNG. Because of that, the ESM would not know if the PW’s state is Active or Standby. As a result, in the downstream direction, traffic could end up being forwarded towards the Standby PW, effectively being black-holed.

To overcome this, SRRP can be used in conjunction with an additional mechanism to help monitor the activity of the PWs. This monitoring mechanism is very similar to Fate-sharing. The difference in this case is that the messaging SAP (instead of SRRP instance) is monitoring the activity of the PW. As a result, the SRRP messaging SAP reflects the state of the PW. For example, the PW in a Standby mode would cause the messaging SAP to be in the DOWN state while the PW Active state would cause the messaging SAP to be in the UP state. That is, the SRRP instance reflects the operational state of the messaging SAP. SRRP is indirectly tied into PW state.

Modifying the priority of SRRP instance based on PW’s state as a mean of mapping the Master SRRP into the Active PW would not help here as SRRP messages are not flowing over standby PWs. This is why SRRP state must be enforced via the messaging SAP.

Fate-sharing for PW termination in conjunction with SRRP is not supported.

Metric adjustment for the subscriber routes is supported. Once the tracked SRRP instance transitions into a non-Master state, the state attribute of the route changes and the appropriate action defined in the routing-policy is taken.

9.3.21.5 Group Monitor

The failure detection mechanism to trigger an action within FSG relies on the operational state of the messaging SAP. Such failure detection mechanism is referred as a group monitor.

Group monitor can also be used to detect the state change of the PW. PW state change is reflected in the messaging SAP which in turn triggers the state change of an SRRP instance.

All this is implemented through an oper-group object which is described in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN*. All entities that needs to be monitored (messaging SAPs and PWs) are associated with this oper-group object. Finally, an SRRP instance (in case of FSG) or a messaging SAP (in case of PW) is instructed to monitor the entities in the oper-group object. State transitions of objects in a oper-group object trigger state transitions of entities that are monitoring them (messaging SAPs and SRRP instances). State transitions of monitored objects in a oper-group will cause the following actions:

- With FSG, priorities of SRRP instances are recalculated
- With PW termination on BNG, the operational state of the messaging SAP is changed.

This is an overview of the CLI syntax showing the principles of how should this work (for exact description of commands and full syntax, see the command reference section):

```
config>service
    oper-group <name>   //oper-group creation

config>service(IES | VPRN)>sub-if>grp-if>sap
    oper-group <name>//adds the SAP to the oper-group
        monitor-oper-group <name> // links the status of the oper-
group to the SAP. In this fashion a messaging SAP can monitor the state of a PW.

config>service(IES | VPRN)>sub-if>grp-if>srrp x
    monitor-oper-group <name> priority-step [0-253] //
with this, a state transition of the objects in the oper-
group should trigger SRRP priority recalculation. The state of the oper-
group is not important but in the state of the objects within. If an object within t
he oper-group goes down, the SRRP priority is lowered by a priority-
step. The SRRP priority will be adjusted on every state transition of member objects
.

config>service>epipe>spoke-sdp
    oper-group <name> // this will add a PW to the oper-
group. A messaging SAP monitors this PW and it assumes the state according to the st
ate of the PW in the oper-
group. A standby or a DOWN PW state causes the messaging SAP to assume a DOWN state.
Otherwise, the messaging SAP would be in the UP state. In order for the SAP to assum
e the DOWN state, both RX and TX side of the PW must be shut. In other words, a PW i
n standby
mode also must have the local TX disabled by the virtue of the ‘slave’ flag (standby
-
signaling-slave command under the spoke-
sdp hierarchy). Without the TX disabled, the
SAP monitoring the PW would not transition in the down state.
```

A hold timer is provided within the **oper-group** command to suppress flapping of the monitored object (SAP or pseudowire).

Example with ESM over pseudowire through a VSM ‘loop’.

Figure 136 **Pseudowire Example**

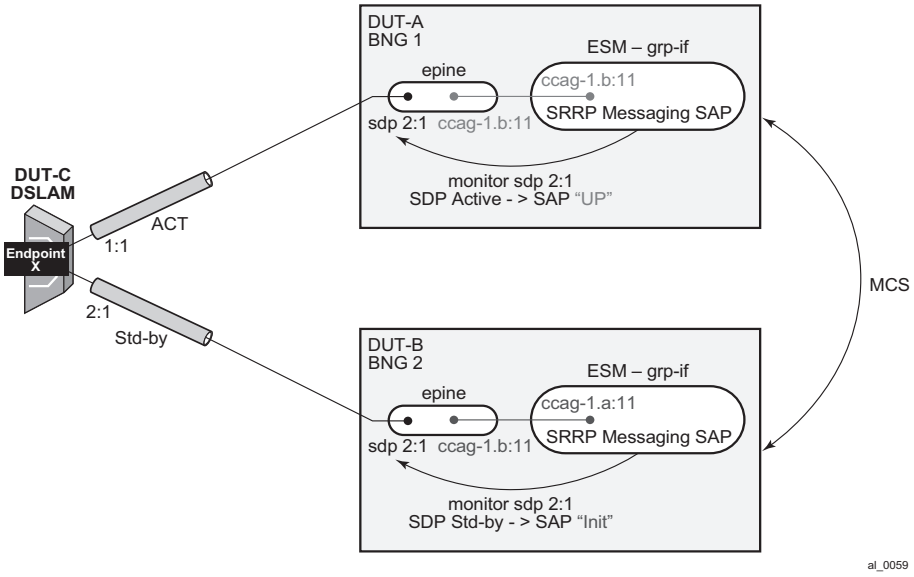


```
*A:Dut-C>config>service>epipe# info
----------------------------------------------
            endpoint "x" create
                standby-signaling-master
            exit
            sap 1/1/7:1 create
            exit
            spoke-sdp 1:1 endpoint "x" create
                precedence primary
                no shutdown
            exit
            spoke-sdp 2:1 endpoint "x" create
                no shutdown
            exit
            no shutdown
----------------------------------------------

*A:Dut-A>config>service>epipe# info
----------------------------------------------
            sap ccag-1.b:11 create
            exit
            spoke-sdp 2:1 create
                standby-signaling-slave
                oper-group "1"
                no shutdown
            exit
            no shutdown
----------------------------------------------
*A:Dut-B>config>service>epipe# info
----------------------------------------------
            sap ccag-1.b:11 create
            exit
            spoke-sdp 2:1 create
```

```
                standby-signaling-slave
                oper-group "1"
                no shutdown
            exit
            no shutdown
-----------------------------------------------
*A:Dut-A>config>service>ies# info
-----------------------------------------------
            redundant-interface "redif11" create
                address 10.1.1.2/24 remote-ip 10.1.1.4
                spoke-sdp 1:1 create
                    no shutdown
                exit
            exit
            subscriber-interface "subif_1" create
                shutdown
                address 10.1.1.2/24 gw-ip-address 10.1.1.100
                group-interface "grpif_1_2" create
                    shutdown
                    redundant-interface "redif11"
                exit
            exit
            subscriber-interface "subTest" create
                address 10.1.1.2/24 gw-ip-address 10.1.1.254
                group-interface "grpTest" create
                    redundant-interface "redif11"
                    sap ccag-1.a:1 create
                    exit
                    sap ccag-1.a:11 create
                        monitor-oper-group "1"
                    exit
                    srrp 11 create
                        message-path ccag-1.a:11
                        no shutdown
                    exit
                exit
            exit
            no shutdown
-----------------------------------------------
*A:Dut-B>config>service>ies# info
-----------------------------------------------
            redundant-interface "redif11" create
                address 10.1.1.4/24 remote-ip 10.1.1.2
                spoke-sdp 1:1 create
                    no shutdown
                exit
            exit
            subscriber-interface "subif_1" create
                shutdown
                address 10.1.1.4/24 gw-ip-address 10.1.1.100
            exit
            subscriber-interface "subTest" create
                address 10.1.1.4/24 gw-ip-address 10.1.1.254
                group-interface "grpTest" create
                    redundant-interface "redif11"
                    sap ccag-1.a:1 create
                    exit
                    sap ccag-1.a:11 create
                        monitor-oper-group "1"
```

```
                    exit
                    srrp 11 create
                        message-path ccag-1.a:11
                        no shutdown
                    exit
                exit
            exit
            no shutdown
-----------------------------------------------
*A:Dut-B>config>service>ies# show srrp
===============================================================================
SRRP Table
===============================================================================
ID        Service        Group Interface                Admin     Oper
-------------------------------------------------------------------------------
11        1              grpTest                        Up        initialize
-------------------------------------------------------------------------------
No. of SRRP Entries: 1
===============================================================================
*A:Dut-A>config>service>ies# show srrp
*A:Dut-A>config>service>ies#
===============================================================================
SRRP Table
===============================================================================
ID        Service        Group Interface                Admin     Oper
-------------------------------------------------------------------------------
11        1              grpTest                        Up        master
-------------------------------------------------------------------------------
No. of SRRP Entries: 1
===============================================================================
```

9.3.22 Subscriber QoS Overrides

Subscriber QoS overrides enable per-subscriber and per-SLA Profile Instance QoS parameter customization to reduce the amount of sub-profiles and sla-profiles that must be configured on the router to cover all needed service level combinations.

Subscriber QoS overrides can be installed at subscriber host or session creation:

- with an Alc-Subscriber-QoS-Override VSA in a RADIUS Access-Accept message
- with a Charging-Rule-Install/Charging-Rule-Definition/QoS-Information AVP in a DIAMETER Gx CCA message

Note: To use the APN-Aggregate-Max-Bitrate-DL and APN-Aggregate-Max-Bitrate-UL AVPs for QoS overrides, a corresponding 3gpp-qos-mapping must be configured in the DIAMETER Gx application policy:

>config>subscr-mgmt>diam-appl-plcy>gx>3gpp-qos-mapping>

[no] apn-ambr-dl - Configure the APN-AMBR mapping for the downlink

[no] apn-ambr-ul - Configure the APN-AMBR mapping for the uplink

Subscriber QoS overrides can be installed, updated or removed in a mid-session change with a RADIUS CoA, a DIAMETER Gx RAR or a DIAMETER Gx CCA message using the same attributes as for a subscriber host or session creation.

Subscriber QoS overrides can also be activated using subscriber services. See QoS Override-based Subscriber Service for details.

The format of the [26.6527.126] Alc-Subscriber-QoS-Override VSA is described in the *7750 SR and VSR RADIUS Attributes Reference Guide*.

The format of QoS Overrides AVP's in the 3GPP-1016 QoS-Information AVP are described in the *7750 SR Gx AVPs Reference Guide*.

The following SLA Profile Instance QoS parameters can be overridden:

- ingress and egress queue: **pir**, **cir**, **mbs**, **cbs**
- ingress and egress policer: **pir**, **cir**, **mbs**, **cbs**
- egress queue class weight (applies to HSQ card only)
- egress queue wrr weight (applies to HSQ card only)
- egress aggregate rate (applies to HSQ card only)
- egress wrr group: rate, class weight (applies to HSQ card only)

The following subscriber QoS parameters can be overridden:

- egress aggregate rate
- ingress and egress root arbiter rate
- ingress and egress intermediate arbiter rate
- ingress and egress user scheduler: **rate**, **cir**
  - The ingress and egress user scheduler overrides through DIAMETER Gx can only be done using APN-Aggregate-Max-Bitrate-UL and APN-Aggregate-Max-Bitrate-DL AVPs and requires the following 3gpp-qos-mapping in the DIAMETER Gx Application policy

```
>config>subscr-mgmt>diam-appl-plcy>gx>
        3gpp-qos-mapping
```

```
            apn-ambr-dl scheduler <scheduler-name>
            apn-ambr-ul scheduler <scheduler-name>
        exit
```

- egress queue: **pir**, **cir**, **mbs**, **cbs**, **wrr weight** (applies to HS-MDA only)

The operational value of some of the QoS parameters can be derived from different sources.

For queue and policer QoS parameters, the following hierarchy applies (highest priority is listed first):

- Credit Control overrides
- Subscriber Services QoS overrides
- Subscriber QoS overrides (RADIUS, DIAMETER)
- Overrides configured at sla-profile level
- Queue parameters set in QoS policy level

For scheduler and arbiter overrides, the following hierarchy applies:

- ANCP overrides
- Subscriber Services QoS overrides
- Subscriber QoS overrides (RADIUS, DIAMETER)
- Overrides configured at sub-profile level
- Scheduler/arbiter parameters as configured in scheduling/policer-control-policy

Up to 18 QoS overrides can be installed per subscriber host or session. A new set of QoS overrides received using a mid-session change replaces the previous set of QoS overrides.

QoS overrides are always stored as part of the subscriber host or session data but are only applied when the override is valid in the active QoS configuration. For example:

- An egress queue 5 PIR rate override is stored with the subscriber session but not applied when the sap-egress QoS policy has no queue 5 defined

RADIUS or DIAMETER Gx initiated QoS overrides can be displayed with the following show commands:

- **show service id** <*service-id* | *name*> **ipoe session detail**
- **show service id** <*service-id* | *name*> **ppp session detail**
- **show service id** <*service-id* | *name*> **dhcp lease-state detail**
- **show service id** <*service-id* | *name*> **dhcp6 lease-state detail**

- **show service id** <*service-id* | *name*> **arp-host detail**
- **show service id** <*service-id* | *name*> **slaac-host detail**

Subscriber services initiated QoS overrides can be displayed with:

- **show service sub-services**

The active QoS overrides per-subscriber and per-SLA Profile Instance can be displayed with:

- **show service active-subscribers detail**

The number of allocated and free Subscriber SLA Profile Instance QoS overrides, QoS Intermediate Arbiter Overrides and QoS User Scheduler Overrides per-line card can be monitored with the **tools dump resource-usage card** CLI command.

Subscriber QoS overrides are synchronized through MCS in a dual-homing environment. QoS overrides are not stored in the subscriber-mgmt application persistence file.

9.3.23 Dual Stack Lite

Dual Stack Lite feature is supported on the 7710 SR-Series in combination with the MS-ISA to function as a DS-Lite Address Family Transition Router (AFTR).

Dual Stack Lite is an IPv6 transition technique that allows tunneling of IPv4 traffic across an IPv6-only network. Dual-stack IPv6 transition strategies allow service providers to offer IPv4 and IPv6 services and save on OPEX by allowing the use of a single IPv6 access network instead of running concurrent IPv6 and IPv4 access networks. Dual-Stack Lite has two components: the client in the customer network, known as the Basic Bridging BroadBand element (B4) and an Address Family Transition Router (AFTR) deployed in the service provider network.

Dual-Stack Lite leverages a network address and port translation (NAPT) function in the service-provider AFTR element to translate traffic tunneled from the private addresses in the home network into public addresses maintained by the service provider. On the 7750 SR, this is facilitated through the Carrier Grade NAT function.

Figure 137 **Dual-Stack Lite**

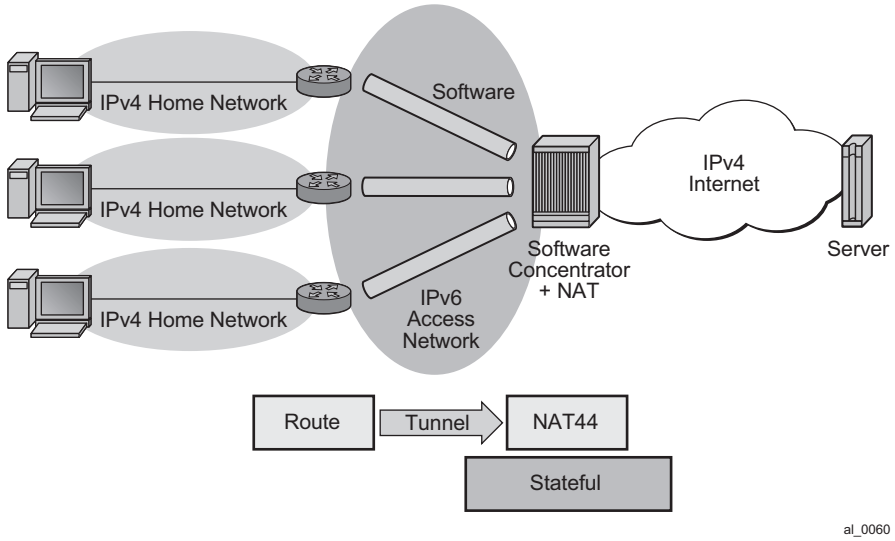

As shown in Figure 137, Dual-Stack Lite has two components, a softwire initiator in the RG and a softwire concentrator, deployed in the service provider network, where control-less IP-in-IP (using protocol 4 - IPv4 in IPv6) is used for tunneling. When using control-less protocol, packets are sent on the wire for the remote softwire endpoint without prior setup of a tunnel.

The softwire initiator in the home network is combined with a routing function, where the default route is passed to the softwire pseudo-interface. Note that there is no NAT function, therefor, the private IP addresses of the home network are encapsulated without source address modification, and forwarded to the softwire concentrator where all NAT is performed. The softwire pseudo-interface unicasts all IPv4 traffic to the IPv6 address of the softwire concentrator, which was pre-configured.

When encapsulated traffic reaches the softwire concentrator, the device treats the source-IP of the tunnel to represent a unique subscriber. The softwire concentrator performs IPv4 network address and port translation on the embedded packet by re-using Large Scale NAT and L2-Aware NAT concepts.

9.3.23.1 IP-in-IP

As shown in Figure 138, IP-in-IP uses IP protocol 4 (IPv4) to encapsulate IPv4 traffic from the home network across an IPv6 access network. The IPv4 traffic tunneling is treated as best-effort with no subscriber management or policy, and does not use ESM. The scale is dependent only on the internal structures of the MS-ISA and CPM, that is, the IP-in-IP model can support more subscribers than an ESM-based approach.

Figure 138 **IP-in-IP**

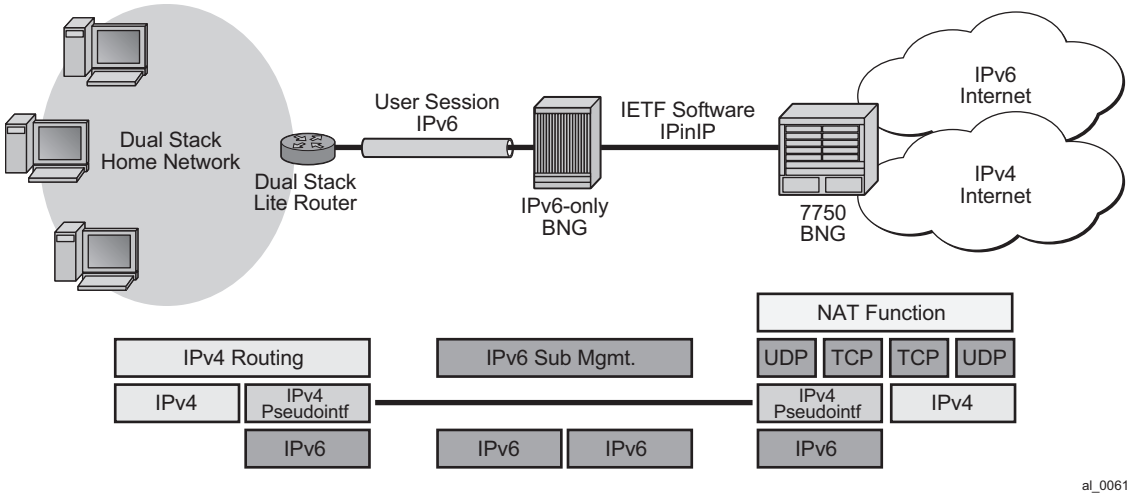

Dual-Stack Lite IP-in-IP is configured through the existing **nat** command that is inside the CLI statements that are within the base router or VPRN. A service performing large scale NAT supports Dual-Stack Lite.

Dual-Stack Lite expects a routing (non-NATing) gateway in the home, where many different IPv4 inside addresses exist for each subscriber. These inside addresses may overlap other subscriber's address, especially given the heavy use of RFC 1918 address space.

The lack of control of protocol for the IP-in-IP tunnels simplifies the functional model, since any received IPv4 packet to the ISA dual-stack-lite address can simply be:

- Checked for protocol 4 in the IPv6 header.
- Checked that the embedded IP packet is IPv4.
- Processed as if it were L2-Aware, where the source-IP of the tunnel (the source IPv6 address) is used as the subscriber identifier.

Note that the inside IP address in the NAT, tables must not be the IPv6 address of the tunnel, but the true IPv4 address of any host within the home. The subscriber-id must be the literal IPv6 address (appreciating this may be 34 characters in length).

9.3.23.2 Configuring Dual Stack Lite

Dual Stack Lite is configured on an inside service and uses the existing Large Scale NAT nat-policies and outside pools. Dual-Stack Lite and NAT44 Large Scale NAT can operate concurrently on the same inside and outside services.

Dual Stack Lite is configured using the following CLI:

```
configure {router | service vprn service-id}
- [no] nat
- inside
- [no] dual-stack-lite
- [no] *address ipv6-address
```

9.3.23.3 L2TP over IPv6

In this mode, L2TP provides the transport for IPv4 that allows full ESM capabilities on the 7750 SR. From the node's perspective, the L2TP tunnel is no different in capability to those already supported. Only the underlying transport (IPv6 instead of IPv4) distinguishes this approach.

To support legacy IPv4 access, L2TP over IPv6 is combined with the existing L2-Aware NAT feature as shown in Figure 139.

As ESM is used, scale is limited by the number of ESM hosts supported on a chassis and any associated resources like queues.

Figure 139 **L2TP over IPv6**

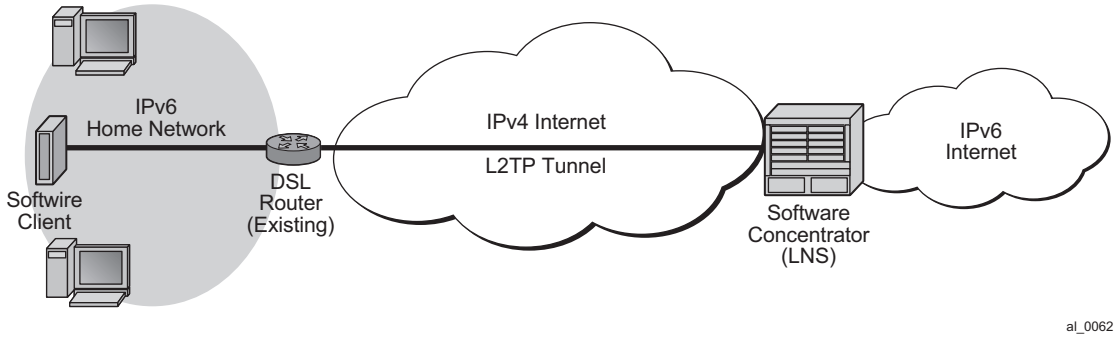

L2TP LNS over IPv6 is supported in both the base routing instance and VPRN that has 6VPE configured.

Like the LNS implementation, tunnels are terminated on any routing interface, including loopback, SAP, or network port. A single interface simultaneously supports IPv4 and IPv6 L2TP tunnel termination by having two different addresses configured.

For greater scalability, L2TP tunnel and session count per chassis are increased to allow 1 tunnel per session.

NAT capabilities are supported via existing L2-Aware NAT methods. Note that the L2TP LNS over IPv6 may be used without NAT as well and the L2TP sessions may be either IPv6-only or dual-stack.

9.3.24 Call Trace

Call trace is an enhanced debugging feature that allows control plane messages for a single entity to be monitored. When call trace is enabled, all protocols related to this entity are captured. Operators can use this information to easily debug problematic entities such as an IPoE session instead of debugging and verifying separate protocols such as DHCP, ARP, or RADIUS.

Call trace will also log some kinds of events that are not directly associated to a protocol, such as LUDB access. These events are encrypted and are intended for research and development purposes.

Call trace can present the captured packets for further processing in one of the following ways:

- Call trace can send the captured packets as live output over a UDP tunnel to an external monitoring device.
- Call trace can store the captured packets as PCAP on a pre-provisioned compact flash. If the system already uses the compact flash intensively, such as for ESM persistency, then this method is not recommended for use in a live network.
- Call trace can display decoded packets in debug output. This method is mainly for use in low-scale debugging of a few entities; use on large-scale live networks may impact overall performance.

Generated traces contain the original packets, encapsulated in a custom header that contains metadata. To decode the metadata and extract the packet, a Nokia-specific Wireshark plug-in is required. Contact the Nokia Technical Assistance Center (TAC) for information.

When storing call trace results on a compact flash, files are not automatically synchronized to the standby CPM.

Call trace distinguishes between traces and trace jobs. A trace consists of a set of matching criteria and additional parameters such as a trace profile and a name. Each session that matches a trace will create a trace job as long as system resources are available. Trace jobs can either be stopped individually or by removing the original trace. By default, existing sessions will not create a trace job when a new trace is enabled; this functionality must be explicitly enabled.

9.4 L2TP Tunnel RADIUS Accounting

Figure 140 **L2TP Tunnel Accounting**

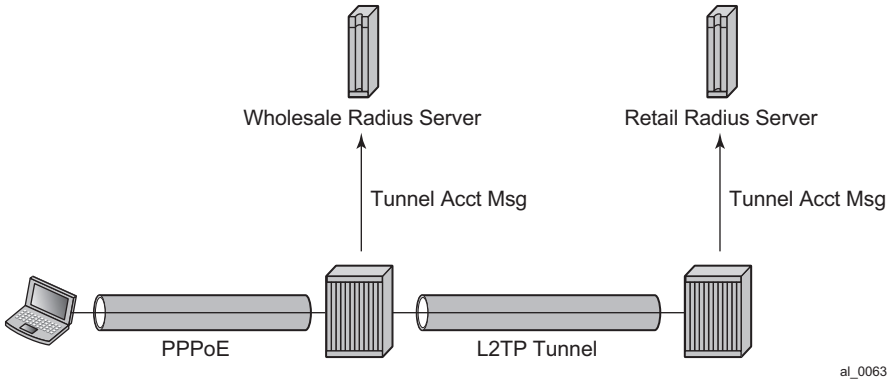

When L2TP tunnel accounting is enabled, except for **host** or **sla-profile**-based accounting packets and attributes, the following are additional accounting packets and attributes:

- Accounting packets: tunnel-start/stop/reject; tunnel-link-start/stop/reject — There are no interim updates for L2TP tunnel/session accounting.
- RADIUS accounting attributes:
  - Tunnel-Assignment-Id (LAC only)
  - Acct-Tunnel-Connection
  - Acct-Tunnel-Packets-Lost

These attributes were added into current account-start/stop/interim-update packets (host accounting/sla-profile accounting)

Tunnel level accounting and session level accounting can be enabled or disabled independently.

New accounting packets and related RADIUS attribute list are described in Table 84.

Some considerations of RADIUS attributes are described in RADIUS Attributes Value Considerations.

9.4.1 Accounting Packets List

Table 84 describes L2TP tunnel accounting behavior along with some key RADIUS attributes (apply for both LAC and LNS):

Table 84 **L2TP Tunnel Accounting Behavior**

Act-Packet	When	Key Attributes	Remark
Tunnel-Start	A new L2TP tunnel is created	Acct-Session-ID	—
		Event-Timestamp	—
		Tunnel-Type:0	—
		Tunnel-Medium-Type:0	—
		Tunnel-Assignment-Id:0	—
		Tunnel-Client-Endpoint:0	—
		Tunnel-Client-Auth-Id:0	—
		Tunnel-Server-Endpoint:0	—
		Tunnel-Server-Auth-Id:0	—
Tunnel-Reject	A new L2TP tunnel creation failed	Acct-Session-Id	—
		Event-Timestamp	—
		Tunnel-Type:0	—
		Tunnel-Medium-Type:0	—
		Tunnel-Assignment-Id:0	—
		Tunnel-Client-Endpoint:0	—
		Tunnel-Client-Auth-Id:0	—
		Tunnel-Server-Endpoint:0	—
		Acct-Terminate-Cause	—
Tunnel-Stop	An established L2TP tunnel is removed	Acct-Session-Id	—
		Event-Timestamp	—
		Tunnel-Type:0	—
		Tunnel-Medium-Type:0	—
		Tunnel-Assignment-Id:0	—

Table 84 **L2TP Tunnel Accounting Behavior (Continued)**

Act-Packet	When	Key Attributes	Remark
		Tunnel-Client-Endpoint:0	—
		Tunnel-Client-Auth-Id:0	—
		Tunnel-Server-Endpoint:0	—
		Tunnel-Server-Auth-Id:0	—
		Acct-Session-Time	—
		Acct-Input-Gigawords	—
		Acct-Input-Octets	—
		Acct-Output-Gigawords	—
		Acct-Output-Octets	—
		Acct-Input-Packets	—
		Acct-Output-Packets	—
		Acct-Terminate-Cause	—
Tunnel-Link-Start	An L2TP session is created	User-Name	—
		Acct-Session-Id	This is the same as Acct-Session-id in access-request of host auth
		Event-Timestamp	—
		Service-Type	Framed
		Class	—
		Tunnel-Type:0	—
		Tunnel-Medium-Type:0	—
		Tunnel-Assignment-Id:0	—
		Tunnel-Client-Endpoint:0	—
		Tunnel-Client-Auth-Id:0	—
		Tunnel-Server-Endpoint:0	—
		Tunnel-Server-Auth-Id:0	—

Table 84 **L2TP Tunnel Accounting Behavior (Continued)**

Act-Packet	When	Key Attributes	Remark
		Acct-Tunnel-Connection	See RADIUS Attributes Value Considerations
Tunnel-Link-Reject	A new L2TP session creation is failed	Acct-Session-Id	Should be as same as Acct-Session-id in access-request of host auth
		Event-Timestamp	—
		Tunnel-Type:0	—
		Tunnel-Medium-Type:0	—
		Tunnel-Assignment-Id:0	—
		Tunnel-Client-Endpoint:0	—
		Tunnel-Client-Auth-Id:0	—
		Tunnel-Server-Endpoint:0	—
		Acct-Terminate-Cause	—
		Acct-Tunnel-Connection	—
Tunnel-Link-Stop	A established L2TP session is removed	User-Name	—
		Acct-Session-Id	Should be as same as Acct-Session-id in access-request of host auth
		Event-Timestamp	—
		Service-Type	Framed
		Class	—
		Tunnel-Type:0	—
		Tunnel-Medium-Type:0	—
		Tunnel-Assignment-Id:0	—
		Tunnel-Client-Endpoint:0	—
		Tunnel-Client-Auth-Id:0	—
		Tunnel-Server-Endpoint:0	—
		Tunnel-Server-Auth-Id:0	—

Table 84 **L2TP Tunnel Accounting Behavior (Continued)**

Act-Packet	When	Key Attributes	Remark
		Acct-Tunnel-Connection	—
		Acct-Session-Time	—
		Acct-Input-Gigawords	—
		Acct-Input-Octets	—
		Acct-Output-Gigawords	—
		Acct-Output-Octets	—
		Acct-Input-Packets	—
		Acct-Output-Packets	—
		Acct-Tunnel-Packets-Lost	—
		Acct-Terminate-Cause	—

Notes:

- Errors will occur if there are multiple hosts sharing the same sla-profile instance and then these hosts go to different tunnel.
- 7750 SRs have an internal limitation of 500 pps for accounting messages. This feature shares the same limitation

9.4.2 RADIUS Attributes Value Considerations

- The value of Acct-Tunnel-Connection uniquely identify a L2TP session, and in order to match LAC and LNS accounting record, the value of Acct-Tunnel-Connection is determined by a method shared by LAC and LNS. This means for a given L2TP session, Acct-Tunnel-Connection from the LAC and LNS are the same.
- Current ESM stats are used in Tunnel-Link and tunnel level accounting. This applies for both standard attribute and the 7750 SR's VSA.
- Tunnel level accounting stats need to aggregate all sessions stats that belong to the tunnel. There may be sessions that traverse before tunnel is down, so the system needs to remember the statistics of every session that has been created within the tunnel.

  This applies for both standard attribute and the 7750 SR's VSA.

- The value of Acct-Tunnel-Packets-Lost is the aggregation of all discarded packets on both ingress and egress.

9.4.3 Other Optional RADIUS Attributes

Table 85 lists the optional attributes that could be optionally included in tunnel accounting packet, some of them are applied for link level accounting only.

Table 85 **Optional RADIUS Attributes**

Attribute	Tunnel/Link
nas-identifier	Both
nas-port	Link level only
nas-port-id	Link level only
nas-port-type	Link level only

9.4.4 RADIUS VSA to Enable L2TP Tunnel Accounting

In order to support pure RADIUS-enabled L2TP tunnel accounting on LAC side, the following RADIUS VSA are supported:

Table 86 **Supported RADIUS VSAs**

VSA	Type	Value
Alc-Tunnel-Acct-Policy	String	Policy-name; if the name is **disable** then this means L2TP tunnel accounting is disabled for this tunnel

The Alc-Tunnel-Acct-Policy takes precedence over what has been defined in CLI when Alc-Tunnel-Group is also returned.

9.4.5 MLPPP on the LNS Side

With MLPPP, the counter on LNS side is only available for the bundle, not for each link, so the SR OS's behavior is:

- For each new link session system sends a tunnel-link-start.
- For each link session that is deleted system sends a tunnel-link-stop.
- For all link sessions except the last one system reports 0 for all counters.
- For the last link session, system reports the actual counters for the bundle.

9.5 RADIUS Route Download

The RADIUS route download mechanism periodically polls a RADIUS server for routes to download. The main objective of this feature is to download, in advance, customer-assigned subnets so that they can be re-advertised to the corresponding routing protocols. In this way, subscriber bring up can potentially be done faster (as the routes are already in place and advertised) and, most importantly, reduce the routing protocol churn as subscribers connect and disconnect. The routes being learned through this mechanism could be both managed routes/delegated prefixes as well as the WAN IP assigned to the subscriber in the case PPPoE and un-numbered interfaces are being used.

The route download process requests the routes to a configured RADIUS server by triggering an access-request message. The key identifier for this message is the username, which is a combination of the system's name (or an optionally configured value), appended by a dash ("-") and then a monotonically increasing integer. The download process sends an access request starting with 1 (such as "hostname-1") and the RADIUS server replies with an access-accept message and a number of routes embedded within the message. The system then increases the counter and sends another access request (this time being hostname-2) and receive a reply with the next batch of routes to download. The process continues, incrementing the counter by 1 each time until the system gets an access-reject or the maximum number of routes that can be downloaded is reached.

The routes to be accepted are in the following format:

[vrf {vprn-name | vprn-service-id}] prefix-mask {null0 | null 0 | black-hole} [metric] [tag tag-value]

The prefix-mask could be in any form as 'prefix/length', 'prefix mask' or 'prefix' (in the latter case, for IPv4 routes, the mask shall be derived from the IP class of the prefix).

The route formats are supported:

- Framed-Route (RADIUS attribute 22)

  Framed-Route = "192.168.3.0 255.255.255.0 null0"

  Framed-Route = "vrf vrfboston 192.168.1.0/24 null 0 0 tag 6"

  Framed-Route = "vrf 2001 192.168.10.0/24 black-hole 0 tag 8"
- Cisco-AVPair (Cisco VSA 26-1)

  cisco-avpair = "ip:route = 192.168.3.0 255.255.255.0 null0"

  cisco-avpair = "ip:route = vrf vrfboston 192.168.1.0/24 null 0 0 tag 6"

IPv6 routes are also supported. The format is based on using the IETF-defined IPv6 Framed-IPv6-Route (attribute 99). The following text shows the supported formats.

- Framed-IPv6-Route (RADIUS attribute 99)

  Framed-Route = "2001:100:bad:cafe::/64 null0"

  Framed-Route = "vrf vrfboston 2100:5aaa:dead:beaf::/96 null 0 0 tag 6"

  Framed-Route = "vrf 3000 2200:1bbbb:dead::/48 black-hole 0 tag 6"

All the routes downloaded will be a new protocol type “**periodic**”. The download process re-starts the AAA requests after a given interval (a configurable value but target refresh rate is 15 minutes) and routes shall be updated according to the following process:

- When the router initiates a new download process, the routes are kept in a temporary table until the download process completes (receives an access-reject from the AAA). The temporary download table is then checked for errors and finally, any changes reflected to the actual routing table.
- Routes no longer present in the download will be removed from the routing table.
- If the AAA server responds with an access-reject for the first username (that is, an implicit empty route-download table), all routes will be removed from the routing table.
- If there are any protocol errors (at the RADIUS level), such as time-out, no response, bad record format, too many records, and so on, the download process is suspended and retried after a configurable timer. The minimum retry timer is at least 1 minute and given the light load this represents control-plane-wise (concurrent downloads are not supported) the retries can continue infinitely until the next refresh period occurs, where the download restarts from the beginning. An exponential backoff algorithm with a configured minimum and maximum delay will be used to determine the retry timer.
- In any case, the routes are only purged from the routing table after a complete download process was achieved (properly terminated with an access-reject message). Under any other failure condition, the routes shall remain active. Shutting down the download process should not remove the downloaded routes. A **clear** command will be provided to clear the periodic routes.
- All the imported routes (blackholes) will be imported into the line-card FIBs to avoid the routing loops caused by announcing the prefixes but not installing the actual blackholes.

9.6 Managed SAP (MSAP)

Managed SAPs allow the use of policies and a SAP template for the creation of a SAP. Although the router supports automatic creation of subscriber hosts in a shared SAP, the most secure mode of operation and common mode is the subscriber per SAP model. In this model, each subscriber is defined with its own VLAN. This feature uses authentication mechanisms supported by the node to provide a SAP.

The reception of a trigger packet initiates a RADIUS or local user database authentication to provide the service context in which the MSAP should be created. The VLAN of the created MSAP is the same as the trigger packet. An MSAP is similar to a regular SAP but its configuration is not user editable and not maintained in the configuration file. The managed SAP remains active as long as the session is active.

The following trigger types are supported:

- DHCP discover (or requests if configured) for DHCP clients. The managed SAP lifetime is defined by the lease time.
- PPPoE PADI for the PPPoE client. The managed SAP lifetime is defined by the session time. The MSAP is installed after the IP address is provided. A short temporary state handles packets between the PADO and ACK.
- ARP defines the managed SAP lifetime. The ARP entry refresh behavior is maintained.
- DHCP6
- PPP
- rtr-solicit (SLAAC hosts)

Multiple trigger types can be enabled on a single capture-sap.

Trigger packets are received on a capture SAP that must be configured in a VPLS service. A capture SAP is defined in a similar way as a regular default SAP but it does not forward traffic. It’s sole purpose is to capture trigger packets for authentication. See the following configuration example:

```
vpls 10 customer 1 create
    sap 1/1/1:*.* capture-sap create
        description "capture sap"
        trigger-packet arp dhcp dhcp6 pppoe
        authentication-policy "auth-policy-1"
    exit
    no shutdown
exit
```

A capture SAP and default SAP cannot be configured at the same time for a single port with the dot1q encapsulation or for a single port:topq combination with qinq encap. The capture SAP is used if a more specific match for the Q or Q-in-Q tags is not found by the IOM. If a capturing SAP is defined, triggering packets are sent to the CPM. Non-triggering packets captured by the capturing SAP are dropped.

An ingress VLAN ID (VID) type mac filter can be configured on a capture-sap to have additional control on the VLANs that are allowed to initiate a host setup. Other filter types are not supported on a capture-sap.

Capture-sap on a dot1q encapsulated port:

- <port-id>:*

  Captures any valid single tagged trigger packet on <port-id> for which no more specific SAP or MSAP is found.

  A single q-tag (<port-id>:tag) is available for authentication.

  The corresponding Managed SAP is created as: <port-id>:tag

Capture-sap on a qinq encapsulated port:

- <port-id>:*.*

  Captures any valid double tagged trigger packet on <port-id> for which no more specific SAP or MSAP is found.

  Both q-tags (<port-id>:tag1.tag2) are available for authentication.

  The corresponding Managed SAP is created as: <port-id>:tag1.tag2.

  The optional **allow-dot1q-msaps** command configured at the capture SAP enables additional support for single tagged trigger packets:

  – Valid single tagged trigger packets for which no more specific SAP or MSAP is found are captured on <port-id>

  – A single q-tag is available for authentication, the second tag is set to zero (<port-id>:tag.0)

  – The corresponding Managed SAP is created as: <port-id>:tag.0

  – The **config>system>ethernet>new-qinq-untagged-sap** command should be configured where a combination of <port-id>:tag1.0 and <port-id>:tag1.tag2 MSAPs co-exist. When not configured, <port-id>:tag1.0 MSAPs attract double tagged <port-id>:tag1.tag2 encapsulated traffic which is either dropped (IPoE traffic) or handled as single tagged traffic causing PPPoE sessions to fail.

- <port-id>:tag1.*

  Captures any valid double tagged trigger packet with outer tag = tag1 on <port-id> and for which no more specific SAP or MSAP is found.

  Both q-tags (<port-id>:tag1.tag2) are available for authentication.

The corresponding Managed SAP is created as: <port-id>:tag1.tag2.

The optional **allow-dot1q-msaps** command configured at the capture SAP enables additional support for single tagged trigger packets:

- Valid single tagged trigger packets with tag=tag1 and for which no more specific SAP or MSAP is found are captured on <port-id>
- A single q-tag is available for authentication, the second tag is set to zero (<port-id>:tag1.0)
- The corresponding Managed SAP is created as: <port-id>:tag1.0
- It is a prerequisite to have the **config>system>ethernet>new-qinq-untagged-sap** command configured to enable both <port-id>:tag1.* capture-sap and <port-id>:tag1.0 MSAP to co-exist. The <port-id>:tag1.0 capture-sap cannot be created when not configured

• <port-id>:*.tag2

Captures any valid double tagged trigger packet with inner tag tag2 on <port-id> for which no more specific SAP or MSAP is found.

Both q-tags (<port-id>:tag1.tag2) are available for authentication.

The corresponding Managed SAP is created as: <port-id>:tag1.tag2.

This is an inverse capture-sap that matches on a fixed inner tag with the outer tag identifying the user. The following restrictions apply when an inverse capture-sap is configured on a port:

- Ethernet ports only
- It is not possible to create y.* saps when there is a *.x capture sap present on the port (y=0,1..4094,* and x=1..4094).
- It is not possible to create a y.* network interface when there is a *.x capture SAP present on the port (y=0,1..4094,* and x=1..4094).
- No support for single tagged managed SAP creation

To enable the creation of single and double tagged managed SAPs via a qinq encapsulated capture SAP, following CLI command should be specified:

```
config service
  vpls 10 customer 1 create
     sap 1/1/1:*.* capture-sap create
        allow-dot1q-msaps
```

In addition, new-qinq-untagged-sap should be configured for scenarios as explained previously:

```
config system
  ethernet
     new-qinq-untagged-sap
```

Note that enabling the **new-qinq-untagged-sap** command affects the behavior of existing <port-id>:tag1.0 SAPs.

Valid single tagged trigger packets will result in the creation of a <port-id>:tag.0 managed sap. With **encap-tag-range** matching in a local user database, it is possible to specify different MSAP defaults for single or double tagged MSAPs. For example:

```
config subscriber-mgmt
        local-user-db "ludb-1" create
            ipoe
                host "single-tagged" create
                    host-identification
                        encap-tag-range start-tag *.0 end-tag *.0
                    exit
                    msap-defaults     # defaults for dot1q MSAPs
                        group-interface "group-int-2"
                        policy "msap-policy-2"
                        service 2000
                    exit
                    no shutdown
                exit
            exit

config service
        vpls 10 customer 1 create
            sap 1/1/1:*.* capture-sap create
                trigger-packet dhcp dhcp6
                allow-dot1q-msaps
                ipoe-session
                    ipoe-session-policy "ipoe-policy-1"
                    user-db "ludb-1"
                    no shutdown
                exit
                msap-defaults       # defaults for qinq MSAPs
                    group-interface "group-int-1"
                    policy "msap-policy-1"
                    service 1000
                exit
            exit
```

A set of mandatory parameters should be provisioned for MSAP creation are as follows:

- **Service id**: The service context in which the MSAP will be created.
- **Interface id**: The name of the group-interface context in which the MSAP will be created. The group-interface must exist in the provided service for the MSAP to be installed (routed CO scenario only).
- **MSAP policy**: The name of the policy that defines the MSAP parameters. The policy must exist in the subscriber-mgmt context.

These parameters can be obtained from the following order of preference:

1. Local user database lookup.
2. RADIUS attributes.
3. Defaults configured at the capture-sap context.

The MSAP parameters can be obtained from a local user database. The local user database should be configured at the capture SAP and group interface context. For example,

IPoEv4 hosts

```
config>service>vpls>sap# dhcp-user-db <local-user-db-name>
config>service>ies>sub-if>grp-if>dhcp# user-db <local-user-db-name>
```

IPoEv6 hosts

```
config>service>vpls>sap# dhcp6-user-db <local-user-db-name>
config>service>ies>sub-if>grp-if>ipv6>dhcp6# user-db <local-user-db-name>
```

PPP hosts

```
config>service>vpls>sap# pppoe-user-db <local-user-db-name>
config>service>ies>sub-if>grp-if>pppoe# user-db <local-user-db-name>
```

When RADIUS authentication is still required after local user database authentication, then the authentication policy must be specified in the local user database. In this case no authentication policy may be configured at the group-interface context. For example,

IPoE hosts

```
config>subscr-mgmt>loc-user-db>ipoe>host# auth-policy <policy-name>
```

PPP hosts

```
config>subscr-mgmt>loc-user-db>ppp>host# auth-policy <policy-name>
```

The MSAP parameters are configured at the local user database host context. For example,

```
config>subscr-mgmt>loc-user-db>ipoe>host# msap-defaults
config>subscr-mgmt>loc-user-db>ppp>host# msap-defaults

  - msap-defaults

 [no] group-interface - Configure the group interface
```

```
 [no] policy          - Configure the MSAP policy
 [no] service         - Configure the service
```

When RADIUS authentication is required to return the MSAP parameters without prior local user database authentication, then the authentication policy should be configured at the capture-sap context. In a Bridged CO model, the authentication policy specified in the capture-sap will also be used for the MSAP in the VPLS service. In a Routed CO model, the same authentication policy must also be configured at the group-interface context. For example,

```
config>service>vpls>sap# authentication-policy <auth-policy-name>
config>service>ies>sub-if>grp-if# authentication-policy <auth-policy-name>
```

The MSAP will not be created if the group-interface name returned from RADIUS has a different authentication policy than the authentication policy configured at the capture-sap.

The following table lists the RADIUS attributes (VSAs) to include in a RADIUS access accept message to obtain MSAP parameters in the RADIUS authentication phase.

Table 87 **RADIUS Attributes**

Attribute Name	Type	Purpose and Format
Alc-MSAP-Serv-Id [26-6527-31]	Integer	Service ID of the service context in which the MSAP will be created.
Alc-MSAP-Policy [26-6527-32]	String	Name of the policy that defines the MSAP parameters.
Alc-MSAP-Interface [26-6527-33]	String	Name of the group-interface context in which the MSAP will be created.

MSAP parameters that are not obtained from a local user database lookup, and that are not returned from RADIUS can be specified in the default-msap section of the capture-sap context (last resort):

```
>config>service>vpls>sap# msap-defaults ?
  - msap-defaults

 [no] group-interface - Configure the group interface
 [no] policy          - Configure the MSAP policy
 [no] service         - Configure the service
```

MSAPs can be created in IES or VPRN group interfaces (Routed CO model) and in a VPLS service (Bridged CO model).

The managed SAP configuration can be persistent. The template MSAP policy is stored with the subscriber host which in turn can be made persistent.

If local user database or RADIUS authentication did not provide all the required information to create the subscriber host (no IP address for example), then the MSAP is created with a short timer while waiting for the host to acquire the missing information. If no host is instantiated when the timer expires, the MSAP is deleted.

Multiple subscribers and/or subscriber hosts can share a single MSAP. The MSAP is created with the first instantiated subscriber host and deleted when the last associated subscriber host is removed from the system. Note that only a single MSAP policy is allowed to be specified for a given MSAP. An attempt to change the MSAP policy by a new subscriber host for an existing MSAP will result in a host setup failure.

MSAPs can be created in a wholesale VPRN service while the corresponding subscriber host or session is terminated in a retail VPRN or IES service. Both wholesale MSAP data (service, group-interface and policy) and retail service id must be provided during authentication.

9.6.1 MSAP QoS Configuration

MSAPs are always used in combination with subscriber management. Subscriber traffic QoS models are defined in policies associated with the SLA- and sub-profiles and result in the instantiation of subscriber queues and policers used for subscriber traffic forwarding. The default QoS policies associated with MSAPs instantiate a single ingress and a single egress queue per MSAP for IES and VPRN services. For VPLS services, an additional ingress multi-point queue is instantiated per MSAP. These MSAP queues have limited use and can be suppressed in most cases. For single-subscriber MSAPs, the MSAP queues can be suppressed with the **sub-sla-mgmt single-sub-parameters profiled-traffic-only** CLI command.

The default QoS policy associated with MSAPs may need to be changed to accommodate different circumstances. For example:

- Saving queue resources when **profiled-traffic-only** cannot be used, such as when more than one subscriber is active on an MSAP

  By mapping all forwarding classes to a policer in the QoS policy associated with an MSAP, a single policer instead of a queue is instantiated on the MSAP.

- Providing adequate QoS treatment for multicast traffic in a per MSAP replication mode

  Egress multicast traffic in per MSAP replication mode is forwarded via the MSAP queues or policers. Multicast traffic can be mapped into a dedicated queue or policer. The MSAP queue can be port-parented to provide scheduling priority on the port level.

The QoS policies associated with an MSAP are configured in the MSAP policy.

9.6.2 Sticky MSAP

After a subscriber session ends, the MSAP is removed from the system and the historical data of the subscriber is deleted. Sticky MSAP allows the MSAP to remain even when the subscriber session ends. This feature is only recommended for service providers who do not oversubscribe MSAPs in the network.

Sticky MSAP provides the following benefits.

- Because the sticky MSAP is never deleted, the subscriber can start a session faster; processing time is reduced because the MSAP does not have to be recreated.
- The MSAP may contain valuable historical information for the service provider. Keeping the MSAP provides a means for the service provider to look up subscriber historical data.

The MSAP is only be eligible for stickiness if it was successfully created. The sticky MSAP introduces a new state: idle. An idle MSAP indicates that the subscriber on the MSAP has disconnected and the MSAP is ready for a new subscriber connection. An example is shown below:

```
A:BNG> # show service sap-using
===============================================================================
Service Access Points
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
 [1/1/20:1841](I)                1000       1     none    1     none   Up   Up
 1/1/1:4000                      1000       1     none    1     none   Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 20
-------------------------------------------------------------------------------
Number of Managed SAPs : 1, indicated by [<sap-id>]
Flags : (I) = Idle MSAP
-------------------------------------------------------------------------------
===============================================================================
```

The **clear service id** *id* **msap** command removes MSAPs.

Note: This command can remove MSAPs with active subscribers. To clear only MSAPs without any active subscriber, use the keyword **idle-only**.

9.6.3 ESM Identification Process

9.6.3.1 SAP-ID ESM Identifier

Providers migrating from Basic Subscriber Management (BSM) can assign a subscriber to a SAP. The SAP ID ESM identifier makes the transition easier by allowing the operator to continue using the *sap-id* as a subscriber-ID.

An ESM SAP ID provides the system the ability to:

- Provide access to the SAP ID string in the Python script.
- Allow the automatic assignment of the SAP-ID to a static subscriber or subscriber host.

9.6.3.2 DSLAM-ID

A DSLAM ID provides the system the ability to define a DSLAM-ID string provided through the Python script, RADIUS, or local user database. If the DSLAM-ID was provided, but the subscriber host is instantiated on a regular MDA (a non-HSMDA), the DSLAM-ID will be ignored.

The HSMDA and the ability to aggregate subscribers into DSLAMs for the purpose of QoS, can use the SAP ID to identify subscribers and associated DSLAMs.

9.6.4 Default-Subscriber

This feature provides a default subscriber definition under the SAP. If the object was configured the operator may use ESM without enabling a processing script or a RADIUS authentication policy. In the event both have been disabled any host that was installed for the SAP will be installed with the configured default subscriber ID. If a RADIUS policy was used or if a script was enabled but a subscriber ID was not returned the default subscriber ID will be used.

9.6.5 Multicast Management

The multicast-management CLI node contains the bandwidth-policy and multicast-info-policy definitions. The bandwidth-policy is used to manage the ingress multicast paths into the switch fabric. The multicast-info-policy is used to define how each multicast channel is handled by the system. The policy may be used by the ingress multicast bandwidth manager, the ECMP path manager and the egress multicast CAC manager.

9.6.6 Subscriber Mirroring

This section describes mirroring based on a subscriber match. Enhanced subscriber management provides the mechanism to associate subscriber hosts with queuing and filtering resources in a shared SAP environment. Mirroring used in subscriber aggregation networks for lawful intercept and debugging is required. With this feature, the mirroring capability allows the match criteria to include a subscriber-id.

Subscriber mirroring provides the ability to create a mirror source with subscriber information as match criteria. Specific subscriber packets can be mirrored mirror when using ESM with a shared SAP without prior knowledge of their IP or MAC addresses and without concern that they may change. The subscriber mirroring decision is more specific than a SAP. If a SAP (or port) is placed in a mirror and a subscriber host of which a mirror was configured is mirrored on that SAP packets matching the subscriber host will be mirrored to the subscriber mirror destination.

The mirroring configuration can be limited to specific forwarding classes used by the subscriber. When a forwarding class (FC) map is placed on the mirror only packets that match the specified FCs are mirrored. A subscriber can be referenced in maximum 2 different mirror-destinations: 1 for ingress and 1 for egress.

9.7 Volume and Time Based Accounting

Time and volume-based accounting includes the following components:

- Metering function performing stateful monitoring of the service delivery to the subscriber.
- Communication with an external management system that gets and updates credit per subscriber, notifications of credit exhaustion, and so on.
- Action on credit exhaustion takes pre-defined action when the credit has been exhausted.

9.7.1 Metering

Metering represents the core of time and volume-based accounting. Service usage is typically measured by performing an accounting of the traffic passing through corresponding subscriber-host queues (volume usage) or by keeping lease-state while the given subscriber-host is connected to the network (time usage).

- Statefullness — The accounting information is compared with pre-defined credit expressed in terms of time or volume to monitor service usage.
- Sensitivity — Defining so called activity-threshold allows distinction between subscriber-host being connected and subscriber-host effectively using the service. This is particularly of interest in cases of time based charging.
- Aggregated usage per-category per-subscriber-host — Accounting information can be reported on per-queue per-sla instance of the given subscriber. In many situations, a certain level of aggregation (such as a per-subscriber or HSI ingress and egress traffic) is required to perform meaningful mechanism for pre-paid services.

9.7.1.1 Categories Map and Categories

This feature introduces a new object category-map which defines individual aggregates (such as data in and out, video and data, and so on) and their mapping to individual forwarding queues.

The following output depicts a category-map configured in the subscriber management context.

```
*A:ALA-48>config>subscr-mgmt# info
----------------------------------------------
```

```
...
        category-map "triple-play" create
            category "data" create
                queue 1 ingress-egress
            exit
            category "video" create
                queue 2 egress-only
            exit
            category "voice" create
                queue 3 ingress-egress
            exit
        exit
        category-map "aggr-subscriber-service" create
            category "data-services" create
                queue 1 ingress-egress
                queue 3 egress-only
            exit
        exit
...
-----------------------------------------------
*A:ALA-48>config>subscr-mgmt#
```

Based on a category-map the system gathers usage information (volume/time) on a per-sla-instance-per-category basis. In order to do so, statistics of all queues and/or policers forming the category of the given sla-instance are aggregated.

- Single subscriber host (routed CPE) — Single SLA instance.
- Multiple subscriber hosts on the same SAP (bridged CPE) — Single SLA instance. Note that several hosts use the same credit and the renewal of one will cause renewal for all.
- Multiple subscriber hosts on different SAP (bridged CPE) — SLA instance per host.

The per-category usage gathered as described above is compared with per-subscriber-host-per-category credit and when credit is exhausted several actions can be taken.

There are several category-maps pre-configured on the system. The category-map applicable to a given subscriber-host will be derived at the host creation from the RADIUS VSA in an authentication-response, Python script, or static configuration in the local-user-database. All subscriber-hosts belonging to the same subscriber and created on the same SAP (hence, sharing the same sla-instance) must use the same category-map. In case of conflict, (an existing subscriber host has a different category-map than the one derived for the new host) the category-map of the last host will be applied to a given sla-instance. As a consequence, all previous information related to the status of the credit will be lost.

There can be multiple queues and/or policers aggregated into one category. There can be up to sixteen categories in a category map.

9.7.1.2 Quota Consumption

There are two types of quota (credit), volume and time. In volume usage monitoring, the system accumulates byte counters per category-sla-instance and compares it with the assigned quota. Once the credit is exhausted (or threshold for renewal is met) the system attempts to renew it with corresponding management system.

In time-based credit, the distinction between active-usage and active-connection is made by defining an activity-threshold, where an object defines an average data rate under which the subscriber-host is considered silent.

As long as the effective rate of the application usage does not exceed the rate defined by the activity-threshold, the given subscriber host will be considered silent and its corresponding credit will not be used. As long as the application usage exceeds the rate, the application-credit will be consumed (in terms of time).

9.7.1.3 Minimum Credit Control Quota Values

The minimum credit control quota values are one second for time quota and one byte for volume quota. These minimum values are not realistic deployment values for multiple reasons such as effective sampling periods, statistics processing time, RADIUS message load, subscriber scale, and so on.

For typical deployment scenarios it is not recommended to implement Credit Control quota values smaller than 60 seconds for time quota and for volume quota the volume that can be consumed in 60 seconds for that category (function of number of queues/policers monitored and their respective rates).

9.7.1.4 RADIUS VSA Alc-Credit-Control-Quota

The quota in the RADIUS VSA Credit-Control-Quota uses this fixed format:

Alc-Credit-Control-Quota = “<volume quota>|<time quota>|<category name>”

- <volume-quota> is specified in bytes (B), in kilobytes (K or KB), in megabytes (M or MB), in gigabytes (G or GB)
- <time-quota> is specified in seconds (s), in minutes (m), in hours (h), in days (d) or a combination (5m30s). A lower unity may never exceed the higher unity: 5m60s is not allowed and should be specified as 6m.

Both volume and time quota should be specified in the attribute but only one credit type (volume or time) is applied per category. The credit-type of a category is configured in the category-map CLI context.

For example, use Alc-Credit-Control-Quota = “0|1h30m|cat1” to grant time quota and Alc-Credit-Control-Quota = “1G|0|cat2” to grant volume quota.

9.7.2 Credit Negotiation Mechanisms

The per-subscriber per-category credit can be obtained by several ways:

- RADIUS during authentication process.
- Static configuration - configured in the **config>subscr-mgmt>category-map>category** context.

Credit can be expressed by either

- Volume
- Time

The renewal of the credit using RADIUS authentication is triggered by credit exhaustion or (if configured) by depletion of the credit to exhausted-credit-threshold level. If this occurs, the system will send a RADIUS authentication message indicating the corresponding category and usage. The following are several possibilities for the RADIUS server response (as shown in Figure 141):

Figure 141 **Threshold Configured/Not Configured**

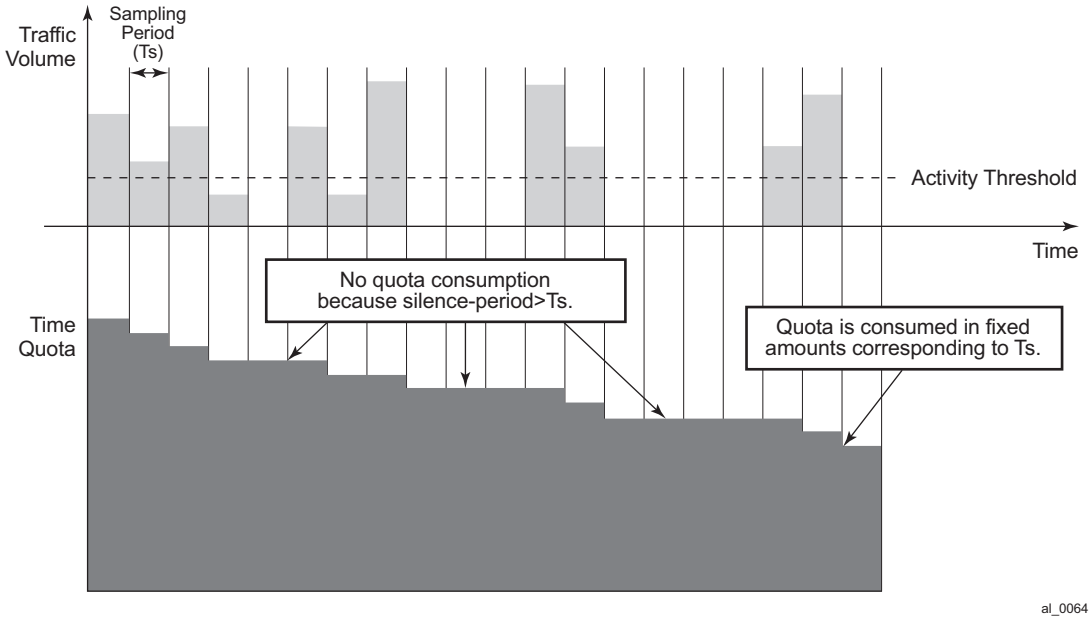

1. No authentication response — The system will install out-of-credit action after the original credit has been used.
2. Authentication response with reject — The corresponding host is removed after the original credit has been used.
3. Authentication response with accept and no credit VSA included — The system will install out-of-credit action.
4. Authentication response with accept and credit VSA included — The credit will be installed.

   The new installed credit is reduced by the amount of credit consumed during credit renewal (i.e. between the start of the credit renewal and reception of an authentication-response). In case the new received credit is less than the credit consumed during credit renewal, then the out-of-credit-action is installed instead.
5. Authentication response with accept and credit VSA included — The out-of-credit will be installed.
   Note that the new credit is always reduced by the amount of credit consumed in time between renewal has been initiated and authentication-respond has been received. In case of a negative result (the newly receive credit is smaller than the amount consumed in the meantime) the test is installed.

In order to identify that the given RADIUS-auth request is related to credit renewal rather than to plain authentication, the node will include empty credit VSAs, depending on categories which has been exhausted. The RADIUS server can identify which category has requested credit renewal.

9.7.3 Action on Credit Exhaustion

System supports configurable actions once the credit for given subscriber is exhausted:

- Sends an SNMP trap and continue (the credit-usage counter is reset).
- Disconnect.
- Changes to a pre-defined service level (such as adjusting the queue rate).
- Blocks the category.

9.7.4 Action on Error-Conditions

During credit negotiation, the number of errors can occur which can lead to a given subscriber-host category with no new credit renewed. This is different from credit exhaustion where a separate configurable action will be taken. The following occurs:

- Sends an SNMP trap and continues.
- Sends a trap and blocks the category.

9.7.5 Applicability of Volume and Time Based Accounting

Volume and time based accounting is applicable to the ESM mode of operation only. Using credit control concept is not mutually exclusive with other accounting methods. In many network implementations the more traditional accounting methods such as XML file or RADIUS accounting will be still used in a combination with the credit concept but with larger intervals. This is helpful when providing overviews of the average usage and service utilization.

9.8 Subscriber Host Idle Timeout

An idle timeout is the maximum time that a subscriber session can be idle before the session is terminated or a connectivity check is started. Idle timeout applies to PPPoE, PPPoEoA, PPPoA and IPoE hosts.

The time/volume based accounting model is used to configure an idle timeout:

- Create a category-map (see Categories Map and Categories)
  - Define a category with queues and/or policers to be monitored for activity (packets being forwarded).
  - An activity threshold (in kb/s) must be configured for idle timeout to take effect. The activity threshold suppresses background traffic (for example control flows) from activity monitoring.

Example:

```
config>subscr-mgmt
        category-map "idle-timeout" create
            activity-threshold 25
            category "cat-1" create
                queue 1 ingress-egress
            exit
        exit
```

In the sla-profile, associate the category-map and optionally define

- An idle-timeout (60 to 15552000 seconds). The default is infinite (no idle-timeout).

  The idle-timeout can also be specified from RADIUS in an access-accept or CoA message with the [28] Idle-Timeout attribute. A RADIUS specified idle-timeout overrides the CLI- configured value. The values outside the limits are accepted but rounded to these boundaries.

Table 88 **Idle-Timeout Attribute**

Attribute ID	Attribute name	Type	Limits	Purpose and Format
28	Idle-Timeout	integer	[60..15552000] seconds	0 = infinite (no idle-timeout) [60..15552000] in seconds For example: Idle-Timeout = 3600

- An idle-action:

– **shcv-check** — Perform a subscriber host connectivity check (IPoE hosts only). Host connectivity verification should be enabled on the corresponding group-interface for the **idle-action shcv-check** to take effect:

config service ies|vprn *service-id* **subscriber-interface** *ip-int-name* **group-interface** *ip-int-name* **host-connectivity-verify**

If the shcv check is successful, the subscriber host is not disconnected and the idle-timeout timer is reset to zero. If the shcv check fails, the subscriber host is disconnected (same as terminate).

For PPP hosts, the **idle-action shcv-check** is ignored and has the same effect as “idle-action terminate”

– Terminate (default): disconnect the subscriber hosts

• IPoE:

Delete the subscriber host

Send a DHCP release message to the DHCP server

Send an Accounting Stop message to the RADIUS accounting server

• PPP:

Delete the subscriber host

Send a terminate request message to the CPE

Send an Accounting Stop message to the RADIUS accounting server

Example

```
config>subscr-mgmt
        sla-profile "sla-profile-1" create
            category-map "idle-timeout"
                category "cat-1" create
                    idle-timeout 3600
                    idle-timeout-action terminate
                exit
            exit
        exit
```

At host instantiation, a timer is initialized to the idle-timeout value (one timer per sla-profile instance). Each queue or policer in the category is monitored for activity over a fixed polling interval:

• During the polling interval:

– if the forwarding rate falls below the configured activity threshold then the timer is deducted by the polling interval (time elapsed).

– If the forwarding rate is above the configured activity threshold then the timer is initialized to the idle-timeout value.

When the timer becomes zero, the idle-timeout-action is performed for all hosts associated with the SLA-profile-instance (all hosts from a subscriber on a single sap and that share the same sla-profile).

9.9 Web Portal Authentication

9.9.1 HTTP-Redirect (Captive Portal)

A captive portal service can be created with an HTTP redirect action in an IP filter. The customer's request to the intended recipient is blocked and the customer is forced to connect to the service's portal server. Refer to HTTP-redirect (Captive Portal) in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Router Configuration Guide* for details.

9.9.2 One-time HTTP Redirection Overview

With this feature enabled, after an ESM host is created, only the FIRST HTTP request from the host will be redirected to a configured URL with specified parameters. Subsequent HTTP request will go through without being redirected.

This feature could be used by service providers to push a web-page to broadband users for purpose of advertisement, announcements, and such.

A **one-time-http-redirection** filter could be configured in **sla-profile**, this filter will be replaced by ingress filter in **sla-profile** after 1st HTTP request is redirected. There is also a RADIUS VSA (Alc-Onetime-Http-Redirection-Filter-Id) that could be included in access-accept or CoA request to override CLI configuration. The format of Alc-Onetime-Http-Redirection-Filter-Id is
Ingr-v4:*filter-id*; for example, **Ingr-v4:1000**. If the filter-id is 0, then system will replace the current **one-time-http-redirection** filter with ingress filter.

In case of CoA, if the host's **one-time-http-filter** has already been replaced then system will just ignore the Alc-Onetime-Http-Redirection-Filter-Id.

If a 7750 SR receives filter insertion via CoA or access-accept when **one-time-http-redirection** filter is still active then the received filter entries will only be applied to the ingress filter. And after 1st http redirection, the update ingress filter will replace the one-time-http-redirection filter.

This feature only supports IPv4 filter.

9.9.3 Web Authentication Protocol (WPP)

The Web Authentication Protocol (WPP) is a protocol running between a BNG and a Web portal server. WPP is used for web portal authentication of WLAN users (DHCP Host). It can function like a web portal that can trigger BNG to perform RADIUS authentication for WLAN users, or send user disconnection notification to BNG.

The Figure 142 illustrates high level of call flow of WPP authentication.

Figure 142 **WPP Authentication**

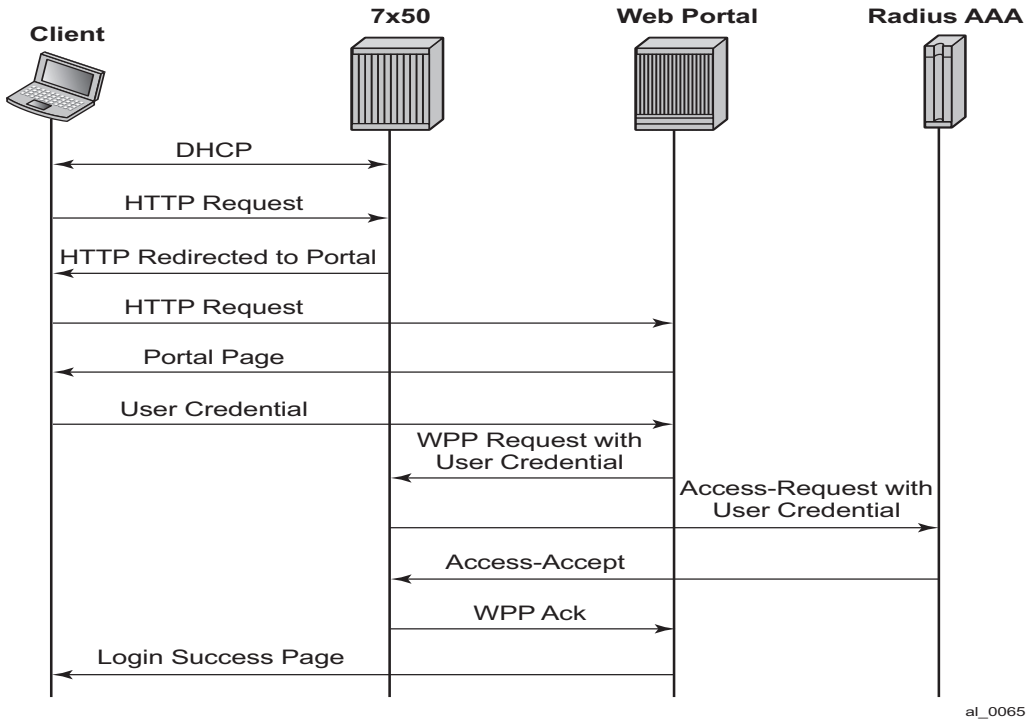

The following describes WPP authentication call flow:

1. When the WLAN user starts a DHCP exchange with a 7750 SR, the router will create a DHCP host from following configurations:
   - Sub-id is the default subscriber ID configured in the **sap>sub-sla-mgmt** context.
   - sla-profile/sub-profile/aa-profile takes the configuration from CLI command **grp-if>wpp>initial-sla-profile/initial-sub-profile/initial-app-profile**.
   - IP address from local or external DHCP server is assigned to the host.
2. When the user sends an HTTP request to visit a website by browser, the router redirects the HTTP request to the web portal.

3. The portal server sends an authentication page to the WLAN user.
4. WLAN user enters username and password in the authentication page and submit to the portal server.
5. The portal server sends a WPP request to router together with the user credentials.
6. The 7750 SR sends an access-request to RADIUS server with user credentials.
7. RADIUS returns an access-accept if authentication succeeds.
8. The 7750 SR returns a WPP ACK to the portal server.
9. If it was access-accept, then the router can optionally override the following host properties:
   - Sub-id: The subscriber ID from RADIUS. If there is no sub-id from RADIUS, then the host keeps using current sub-id.
   - Sla-profile, sub-profile, or aa-profile: The system will use the RADIUS server returned values. If the RADIUS server did not return these then the system will try to use the LUDB (in local DHCP server) return values if they are available. If not, the system will try to use the default values configured under SAP.

9.9.3.1 WPP Configurations

A minimal WPP configurations must include the following:

- WPP portal server — Specifies the name and IP address of the WPP portal server.
- Enable WPP under the group-interface:
  - WPP portal server that system should listen to.
  - **authentication-policy** on **group-interface** that specifies address of RADIUS server.
  - **def-sub-id** under **sap>sub-sla-mgmt** that is used for DHCP host before user is authenticated by portal server.
  - **initial-sla-profile** and **initial-sub-profile** that are used for the DHCP host before user is authenticated by portal server.
    The **initial-sla-profile** should include a ingress filter that has **http-redirection** entry.

The following is an example configuration:

```
#--------------------------------------------------
echo "Web Portal Protocol Configuration"
#--------------------------------------------------
        wpp
```

```
            portals
                portal "portal-1" address 10.9.9.9 create
                    no shutdown
                exit
            exit
            no shutdown
        exit
config>service>vprn# info
----------------------------------------------

            subscriber-interface "sub-if" create
                address 192.168.10.1/24
                group-interface "grp-if" create
                    dhcp
                        server 10.1.1.1
                        gi-address 192.168.10.1
                        no shutdown
                    exit
                    authentication-policy "radius-auth"
                    sap 1/1/9 create
                        sub-sla-mgmt
                            def-sub-id "WLAN-User-Unauth"
                            no shutdown
                        exit
                    wpp
                        initial-sla-profile "webportal"
                        initial-sub-profile "webportal"
                        portal router "Base" name "portal-1"
                        no shutdown
                    exit
                exit
            exit
…
----------------------------------------------
```

9.9.3.2 WPP Triggered Host Creation

In some cases, a 7750 SR can sit behind a Layer 3 device (such as an CMTS), where the router does not participate in client's DHCP process. Such a use case is different from a normal WPP use case where the routers rely on getting client's DHCP request to create an initial ESM host.

This feature allows the system to create an ESM host upon successful WPP authentication without creating an initial host.

In the above use case (behind a Layer 3 device) the user also need to configure one or more default hosts on the SAP to allow HTTP redirection without an ESM host. The default-host subnet is the user's source subnet and the nexthop address is the Layer 3 device's interface address that connect to the SAP. Users also need to configure the **lease-populate l2-header** command in the **grp-if>dhcp** context to make HTTP redirection with default-host work. The **grp-if>dhcp** context could be shutdown in the meantime.

This feature does not work with wholesale/retail

9.9.3.2.1 LUDB Support For WPP

The SR OS supports LUDB lookup for WPP authentication. Users can optionally configure LUDB using the **grp-if>wpp** context to return the WPP-related configuration attributes (such as a portal name, initial-sla-profile, initial-sub-profile, and so on) for an IPoE host. The system can access LUDB when creating the initial host before WPP authentication. The LUDB returned attribute overrides the corresponding configuration under the **group-interface** context.

A LUDB lookup is performed by the system in the following cases.

- when a host is created
- when the system restores a host after a disconnect

If the WPP LUDB lookup returns an authentication policy, it will be used for WPP RADIUS authentication. When WPP LUDB is configured, the authentication policy on **group-interface** is optional and only used by the WPP if there is no authentication policy returned from the WPP LUDB lookup.

9.9.3.3 WPP Multi-Chassis Redundancy Support

The SR OS supports multi-chassis redundancy to WPP. This can be achieved by doing following:

- Create a loopback interface on both 7750 SRs with the same IP address X.
- Use the **track-srrp** parameter while configuring address X to track the corresponding SRRP instance.
- Configure a portal with the same name and same service-id on both nodes to send WPP packets to the destination address.
- Use an route-policy to export X to the routing protocol. The metric the route X can be set is based on the specified SRRP state (master or non-master) so that master node can advertise route X with a better metric. Then the WPP packet from the portal will be attracted to the master.
- Only the master process WPP packet, however in case of standby node receives (such as routing is still re-converging) the WPP packet, then standby will shunt the WPP back to master.
- WPP hosts will be synced via MCS.

9.9.3.4 WPP Portal Group

A WPP portal group allows users to configure up to 8 WPP portals in a portal group. The system can receive portal-initiated WPP request packets from any configured portal in the portal group. When the system must initiate a WPP NTF_LOGOUT message, it sends a NTF_LOGOUT message to all configured portals in the portal group, and the first received ACK_LOGOUT stops retransmission of the NTF_LOGOUT message.

A WPP portal group can be used to achieve WPP portal redundancy:

- Each portal is only allowed to be configured in a single portal group
- Each is allowed to be in a portal group and be used as an individual portal simultaneously
- Mixed WPP versions (version 1 and version 2) of portals are allowed in the same portal group

This feature is also supported for WPP triggered hosts and SRRP/MCS.

9.10 ESM over MPLS Pseudowires

This feature allows IPoE and PPPoE (terminated or L2TP tunneled) subscriber sessions to be backhauled through an Ethernet aggregation network using MPLS pseudowires terminating directly on the BNG. The MPLS pseudowire originates from the first hop aggregation PE (referred to as access PE) upstream of the Access-Node (or directly from a multi-service AN), and terminates on the BNG. Multiple subscriber sessions from a given access-port on the Access-PE can be backhauled over a single P2P MPLS pseudowire towards the BNG. This capability allows the network to scale and does not require a MPLS pseudowire per subscriber between Access-PE and the BNG. The access-port on the Access-PE can be dot1q, q-in-q or NULL encapsulated. The BNG terminates the MPLS pseudowire, decapsulates the received frames, and provides ESM functions including HQoS, without requiring an internal or external loopback. Each MPLS pseudowire is represented on the BNG as a “PW-port” for which SAPs are created. A PW-port can be configured with capture SAP. Both static and managed SAPs are supported. The underlying Ethernet port is required to be in hybrid mode. The feature set is supported for IOM3-XP and HSMDAv2. This feature is supported on the 7750 SR and 7450 ESS in mixed mode.

Figure 143 **ESM over MPLS Pseudowire Example**

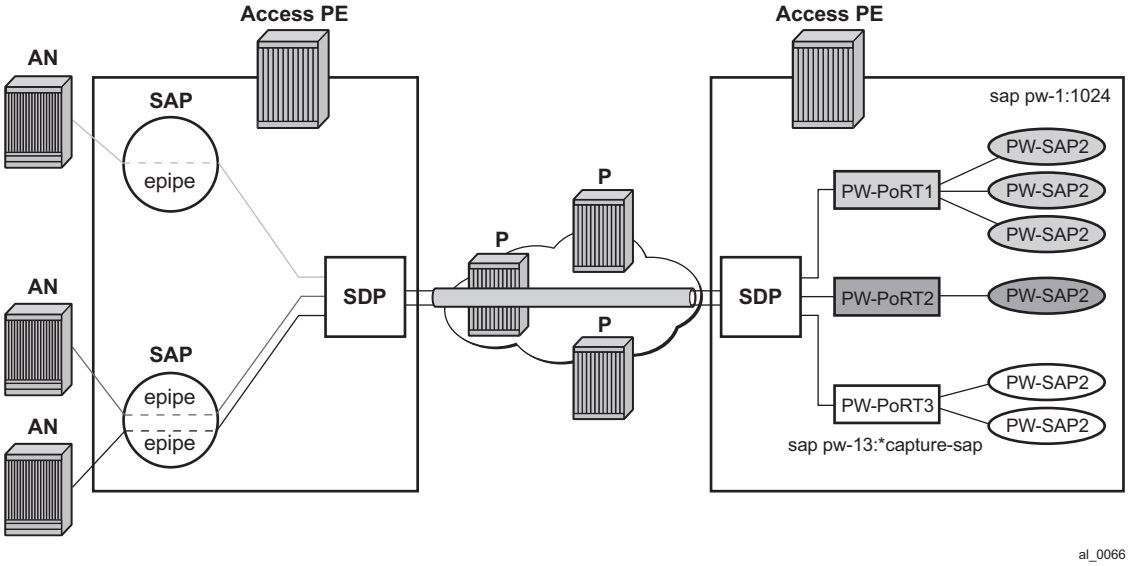

Figure 144 **Group Interface Example**

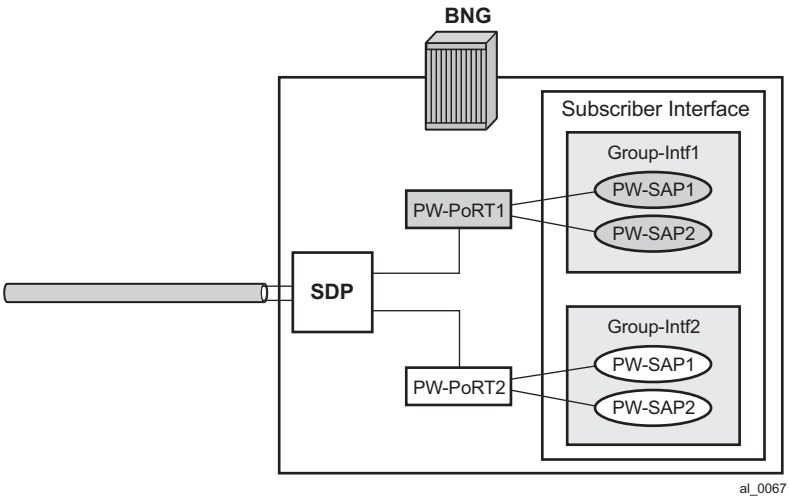

9.10.1 Encapsulation

The subscriber frame encapsulated within the pseudowire is shown in Figure 145. Optional control word is not supported. The SDP could be LDP, RSVP or LDP over RSVP. Hash labels are not supported. SDP is bound to a port or a LAG. In case the SDP is re-routed, the corresponding PW-ports are operationally brought down. The PW-ports are associated with the SDP by configuration.

Figure 145 **Subscriber Frame with PWE Encapsulation**

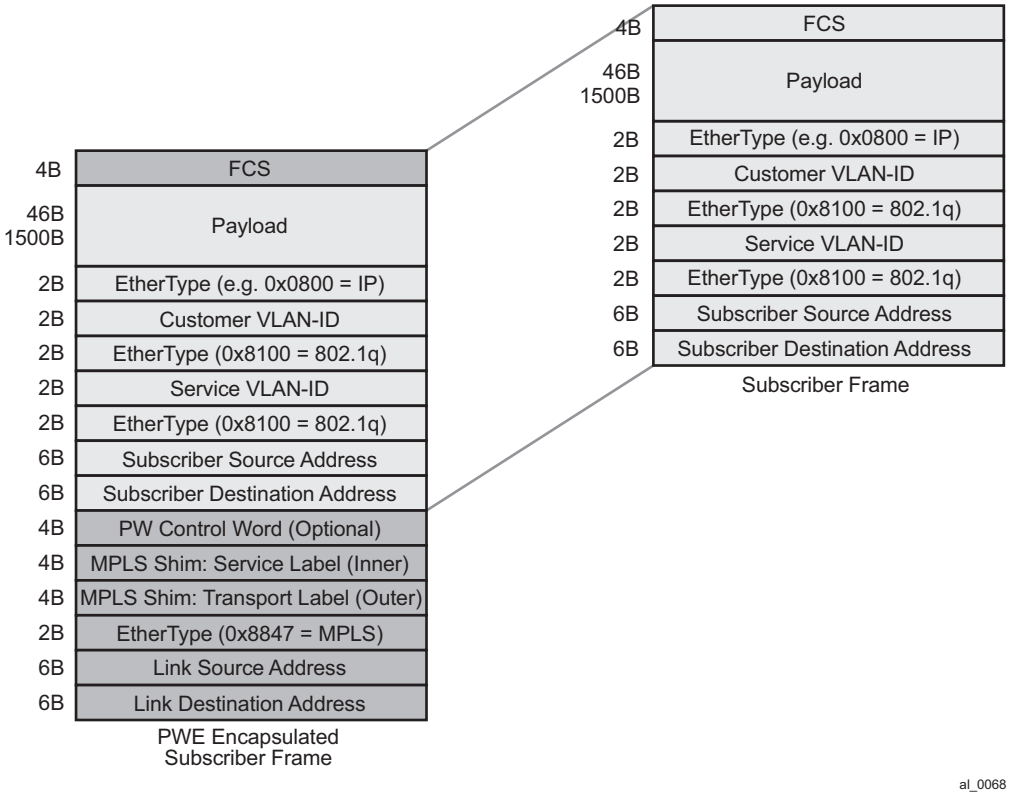

9.10.2 ESM and PW Ports

In BNG deployments, PWs are typically terminated on PW-ports where payload is extracted and processed by ESM on PW-SAPs. There are two modes of operation for PW-ports in SR OS:

- When a PW-port is bound to a specific physical port. A successful mapping between the PW and PW-port requires that the PW terminates on the same physical port (I/O port) to which the PW-port is bound. In this mode of operation, PW-ports do not support re-routing of PWs between the I/O ports. For example, if a PW is rerouted to an alternate physical port due to a network failure, the PW-port will become non-operational.

- When a PW-port is independent of the physical port (I/O ports) on which the PW is terminated. The PXC based PW-port is anchored/bound on the termination side of a port cross connect PXC (which can be either a pxc sub-port or a pxc lag) under an internally created fwd-path-ext (FPE) PW. The benefit of peer 10.20.1.2 create such PW-Port is that it continues to provide services in cases where PW is switched between the I/O ports due to a network failure.

9.10.2.1 ESM on PW-Port Bound to a Physical Port

9.10.2.1.1 QoS Support

QoS is supported for ESM over PW SAPs as with ESM over regular SAPs, and includes currently supported models.

- FC to queue mapping
- H-QoS
  - Per-subscriber H-QoS (service scheduler child to port-scheduler parent).
  - PW SAP queues attached to H-QoS scheduler by a **parent** statement.
  - Scheduler attached to port scheduler by a **port-parent** statement.
- Direct service queue to port-scheduler.
  - Aggregate-rate-limit.

Bandwidth Control at PW Port Level via Vport

Bandwidth control per PW-port (per AN or per AN/ per service), via Vport.

- The Vport can be created on the binding port.
- The Vport can be associated with the PW port either via static assignment or dynamic selection via inter-dest-id (returned from RADIUS or DHCP for a host).
- Aggregate-rate-limit can be configured to shape the egress traffic across all hosts associated with the Vport via inter-dest-sting match or via static association of underlying PW port with the Vport.

The following output displays a dynamic Vport selection based on an inter-dest-id configuration.

```
config>
   Port 1/1/1
      ethernet
         mode hybrid
```

```
         encap-type dot1Q
         mtu 1540
         access
            egress
            vport "v1" create
               agg-rate-limit 1000
               host-match dest "dslam-1"    #### hosts will be associated with
            exit                            #### vport based on inter-dest-id
          exit
       exit
   exit

config>service>sdp>binding
   pw-port 11 vc-id 11 create
       egress
          shaping int-dest-id "dslam-1"   #### dynamic vport selection based on
                                          #### int-dest-id.
```

The following output displays a static assignment of PW port to Vport configuration.

```
config>
   Port 1/1/2
      ethernet
         mode hybrid
         encap-type dot1Q
         access
            egress
            vport "v2" create
               agg-rate-limit 1000
            exit
          exit
      exit
   exit

config>service>sdp>binding
    pw-port 20 vc-id 20 create
       egress
          shaping vport "v2"    #### static assignment of pw-port to vport.
       exit
    exit
```

Last Mile Shaping

With normal Ethernet aggregation in the next-mile, when last-mile shaping is on, fixed encapsulation-offset is calculate based on the last-mile encapsulation type and the next-mile encapsulation (26 Bytes with q-in-q). This offset is applied to the frame, and the ATM overhead is then dynamically calculated on the adjusted size. The resulting dynamically calculated overhead in the data-path is then applied to the queue-rates and the subscriber aggregate-rate.

With this feature of backhauling subscriber sessions using MPLS PW in the aggregation network. The last mile does not see any MPLS PW overhead. The next-mile includes overhead due to the PW encapsulation shown in Figure 145. Therefore, when last mile shaping is enabled, the fixed encapsulation-offset is calculated based on the difference between last-mile encapsulation type and next-mile encapsulation, The next-mile encapsulation takes into account the additional PW overhead, which includes:

14B Ethernet header + [4B] (optional network interface Q-tag) + MPLS Labels (variable)

In the data-path the actual PW encapsulation overhead, taking into account the MPLS labels which could be variable (with FRR or PHP) is tracked, and is applied to the computed “encapsulation offset”. This adjusted “encapsulation offset” is applied to the frame. The ATM overhead is then dynamically calculated on the adjusted size, and applied for last mile shaping (to queue-rates and subscriber-aggregate-rate). Note that there is no change from ESM over normal SAPs, in how last-mile shaping is triggered or how the last mile encapsulation type is determined (via configuration in egress context of subscriber profile or dynamically learned from Access-Loop-Encapsulation sub-TLV in vendor specific PPPoE tags).

9.10.2.1.2 BNG Redundancy with ESM over Pseudowire

This feature provides support for stateful BNG redundancy (when the far-end aggregation PE (A-PE) is dual-homed to two BNGs terminating subscriber sessions over MPLS pseudowires (PWs) that are initiated from the A-PE and provides ESM). Subscriber state between BNGs is synced using MCS.

Epipe-Based Aggregation Service

For an Epipe based aggregation service, the redundancy is based on active/standby PWs from A-PE to dual BNGs. A-PE signals active/standby pseudowire status to peer BNGs. An SRRP instance per PW port (group-interface) is required on the BNG with messaging SAP on each PW port. BNG terminating active PW assumes the mastership for the SRRP instance on the corresponding PW port. SRRP state is tied to the state of the messaging SAP. The messaging SAP goes down when the underlying PW-port goes down, based on PW status bit signaled by the A-PE.

In this model, there is no SRRP message exchange between the two BNGs, as there is no Layer 2 path between the BNGs. The purpose of SRRP is to get SRRP-aware routing for subscriber routes and managed routes, and/or to be able to use the redundant (shunt) interface. Downstream traffic for a subscriber that ingresses the backup BNG can only be shunted to the active BNG, if the corresponding subscriber-interface on the backup BNG is operationally UP. This can be achieved by creating a second empty group-interface (without SAPs) on the same subscriber-interface with the knob 'oper-up-while-empty' configured. Multiple PWs with endpoint configuration is not supported on the BNG.

Figure 146 **BNG Redundancy Based on Active/Standby PW Signaling**

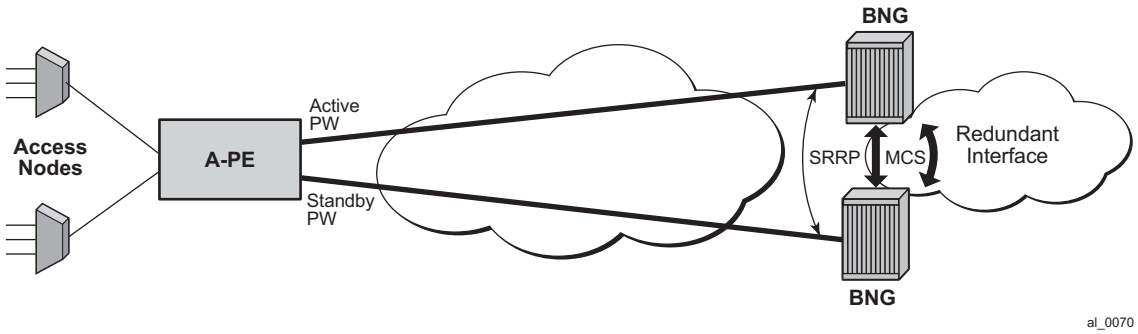

Sample Configuration on Master BNG

```
config>
    pw-port 2 create
    exit

config>redundancy#
     multi-chassis
         peer 10.20.1.3 create
             source-address 10.20.1.2
                sync
                    srrp
                    sub-mgmt ipoe pppoe
                    port pw-2 sync-tag "tag2" create
                    exit
                    no shutdown
                exit
                no shutdown
            exit
        exit
     exit

config>service>ies#
       redundant-interface "redundant-interface" create
          address 10.10.30.2/24 remote-ip 10.10.30.3
          spoke-sdp 23:1000 create
             no shutdown
          exit
```

```
       exit
config>service#
    sdp 1 mpls create
       far-end 10.20.1.2
       ldp
       keep-alive
          shutdown
       exit
       binding
          port 1/1/1
          pw-port 2 vc-id 2 create
             vc-type vlan        #### default encaps-type dot1Q
             no shutdown
          exit
       exit
       no shutdown
    exit
config>service#
       subscriber-interface "subif" create
           address 10.11.1.2/16 gw-ip-address 10.11.1.1 populate-host-routes
                group-interface "grpif" create
                    authentication-policy "base_authpolicy"
                    redundant-interface "redundant-interface"
                    sap pw-2:1000 create
                        description "sap-grp-3"
                    exit
                    srrp 1 create
                        message-path pw-2:1000
                        no shutdown
                    exit
                    arp-host
                        host-limit overall 8000
                        min-auth-interval 1
                        no shutdown
                    exit
                exit
            exit
        exit
```

Sample Configuration on Slave BNG

```
config>
    pw-port 2 create
    exit
config>redundancy#
     multi-chassis
         peer 10.20.1.2 create
             source-address 10.20.1.3
                sync
                    srrp
                    sub-mgmt ipoe pppoe
                    port pw-2 sync-tag "tag2" create
                    exit
                exit
                no shutdown
            exit
        exit
     exit
```

```
config>service>ies#
       redundant-interface "redundant-interface" create
          address 10.10.30.3/24 remote-ip 10.10.30.2
          spoke-sdp 32:1000 create
             no shutdown
          exit
       exit
config>service#
    sdp 1 mpls create
       far-end 10.20.1.2
       ldp
       keep-alive
          shutdown
       exit
       binding
          port 1/1/1
          pw-port 2 vc-id 2 create
             vc-type vlan        #### default encaps-type dot1Q
             no shutdown
          exit
       exit
       no shutdown
    exit
config>service#
       subscriber-interface "subif" create
           address 10.11.1.3/16 gw-ip-address 10.11.1.1 populate-host-routes
                group-interface "grpif" create
                    authentication-policy "base_authpolicy"
                    redundant-interface "redundant-interface"
                    sap pw-2:1000 create
                        description "sap-grp-3"
                    exit
                    srrp 1 create
                        keep-alive-interval 1
                        message-path pw-2:1000
                        no shutdown
                    exit
                    arp-host
                        host-limit 8000
                        min-auth-interval 1
                        no shutdown
                    exit
                exit
                group-interface "dummy" create
                    oper-up-while-empty
                exit
            exit
        exit
```

Sample Configuration on A-PE

```
config>service>epipe#
     description "Default epipe description for service id 103"
     service-mtu 1492
     service-name "XYZ Epipe 103"
     endpoint "x" create
```

```
    standby-signaling-master
exit
sap 1/1/3 create
     description "Default sap description for service id 103"
exit
spoke-sdp 1:2 endpoint "x" create
     description "Description for Sdp Bind 1 for Svc ID 103"
     precedence primary
     no shutdown
exit
spoke-sdp 2:2 endpoint "x" create
     description "Description for Sdp Bind 2 for Svc ID 103"
     no shutdown
exit
no shutdown
```

VPLS-Based Aggregation Service

With VPLS based aggregation service from A-PE, normal SRRP message exchange can take place between the primary and backup BNGs. Master-ship decision and switch-over is based on SRRP. SRRP instance is configured per group-interface corresponding to PW port. Fate-sharing groups (FSG) can be configured for a set of SRRP instances (for example, SRRP instances corresponding to PW ports sharing the same subnet). Standard **oper-group** *grp-id* would need to be configured with messaging SAPs for all PW ports that are in the same FSG, and **monitor-oper-group** *grp-id* would need to be configured under each SRRP instance in same FSG. Existing SRRP support defined in Triple-play services guide for ESM over regular group-interfaces and subscriber SAPs is applicable identically to ESM over PW ports and PW SAPs.

With PW over ESM, redundancy in the aggregation network based on MC-LAG between A-PE and dual BNGs is not supported.

Figure 147 **BNG Redundancy with VPLS Based Aggregation Service**

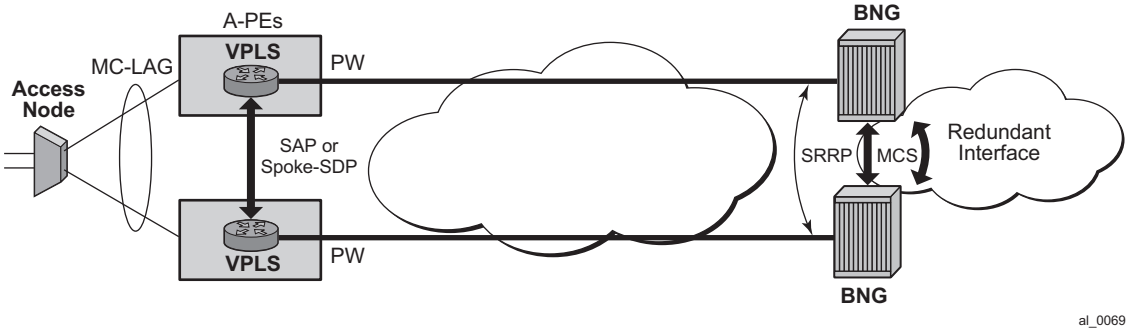

Sample BNG Redundancy Configuration with VPLS Service on A-PE

```
config>
    pw-port 1 create
    exit

config>redundancy#
     multi-chassis
         peer 10.20.1.2 create
             source-address 10.20.1.3
                sync
                    srrp
                    sub-mgmt ipoe pppoe
                    port pw-1 sync-tag "tag1" create
                    exit
                exit
                no shutdown
            exit
        exit
     exit

 config>service>ies
       redundant-interface "red-1-1" create
           address 10.1.1.2/24 remote-ip 10.1.1.1
               spoke-sdp 1:1 create
                  no shutdown
               exit
           exit

       subscriber-interface "sub-1-1" create
            address 10.1.2.2/16 gw-ip-address 10.1.255.254 track-srrp 1
            address 10.2.2.2/16 gw-ip-address 10.2.255.254 track-srrp 2
            dhcp
                gi-address 10.1.2.2
            exit
            group-interface "grp-1-1-1" create
                 srrp-enabled-routing
                 arp-populate
                 dhcp
                     server 10.20.1.2
                     trusted
                     lease-populate 32767
                     client-applications dhcp ppp
                     gi-address 10.1.2.2
                     no shutdown
                 exit
                    authentication-policy "iesAuthPol"
                    redundant-interface "red-1-1"

                  sap pw-1:1.1 create
                      sub-sla-mgmt
                          def-sub-profile “sub_prof_1”
                          def-sla-profile “sla_prof_1”
                          no shutdown
                  exit
                  sap pw-1:4000.1 create
```

```
                    oper-group "1"
                exit

                srrp 1 create
                    gw-mac 00:00:5e:00:01:01
                    keep-alive-interval 50
                    message-path pw-1:4000.1
                    monitor-oper-group "1" priority-step 10
                    no shutdown
                exit
            exit
```

Sample BNG Redundancy Configuration with VPLS Service on A-PE

```
config>service
     customer 1 create
         description "Default customer"
     exit
     sdp 1000 mpls create
         far-end 10.20.1.2
         lsp "lsp_1"
         path-mtu 1600
         keep-alive
         no shutdown
      exit
      sdp 1002 mpls create
         far-end 10.20.1.3
         lsp "lsp_3"
         path-mtu 1600
         keep-alive
         no shutdown
      exit
      vpls 1 customer 1 create
          service-mtu 1600
          stp
          sap 1/1/2 create   // to Access-Node
          exit
        sap 1/1/3 create; //to A-PE2
        exit
          spoke-sdp 1000:1 create  // to BNG1
              no shutdown
          exit
          no shutdown

        exit
    exit
```

A-PE Configuration with VPLS Aggregation Service (A-PE2)

```
config>service
     customer 1 create
         description "Default customer"
     exit
```

```
 sdp 1002 mpls create
     far-end 10.20.1.3
     lsp "lsp_2"
     path-mtu 1600
     keep-alive
     no shutdown
  exit

  vpls 1 customer 1 create
      service-mtu 1600
      stp
      sap 1/1/2 create   // to Access-Node
      exit
    sap 1/1/3 create; //to A-PE1
    exit
      spoke-sdp 1002:1 create  // to BNG2
          no shutdown
      exit
      no shutdown
    exit
exit
```

Show Commands Related to Active/Standby Pseudowire on Dual BNGs

The following example shows SRRP status, subscriber host, and routing information on master BNG:

```
A:Dut-B>config>redundancy# show srrp 1
===============================================================================
SRRP Instance 1
===============================================================================
Description         : (Not Specified)
Admin State         : Up                Oper State       : master
Preempt             : yes               One GARP per SAP : no
Monitor Oper Group  : None
System IP           : 10.20.1.2
Service ID          : VPRN 3
Group If            : grpif             MAC Address      : 1c:85:ff:00:00:00
Grp If Description  : N/A
Grp If Admin State  : Up                Grp If Oper State: Up
Subscriber If       : subif
Sub If Admin State  : Up                Sub If Oper State: Up
Address             : 10.11.1.2/16      Gateway IP       : 10.11.1.1
Redundant If        : redundant-interfa*
Red If Admin State  : Up                Red If Oper State: Up
Address             : 10.10.30.2/24
Red Spoke-sdp       : 23:1000
Msg Path SAP        : pw-2:1000
Admin Gateway MAC   :                   Oper Gateway MAC : 00:00:5e:00:01:01
Config Priority     : 100               In-use Priority  : 100
Master Priority     : 100
Keep-alive Interval : 1 deci-seconds    Master Since     : 05/29/2012 07:22:26
Fib Population Mode : all
VRRP Policy 1       : None              VRRP Policy 2    : None
===============================================================================
```

```
* indicates that the corresponding row element may have been truncated.

A:Dut-B>config>redundancy# show service id 3 arp-host
===============================================================================
ARP host table, service 3
===============================================================================
IP Address      Mac Address       Sap Id              Remaining           MC
                                                      Time                Stdby
-------------------------------------------------------------------------------
10.11.1.11      00:80:00:00:00:01 [pw-2:11]           03h35m47s
10.11.1.12      00:80:00:00:00:02 [pw-2:12]           03h35m47s
-------------------------------------------------------------------------------
Number of ARP hosts : 2
===============================================================================

A:Dut-B>config>redundancy# show router 3 route-table 10.11.1.11
===============================================================================
Route Table (Service: 3)
===============================================================================
Dest Prefix[Flags]                              Type    Proto    Age         Pref
      Next Hop[Interface Name]                                     Metric
-------------------------------------------------------------------------------
10.11.1.11/32                                   Remote  Sub Mgmt 00h24m26s   0
       [grpif]                                                        0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: L = LFA nexthop available    B = BGP backup route available
       n = Number of times nexthop is repeated
===============================================================================
A:Dut-B>config>service>vprn#
```

The following shows SRRP status, subscriber host, and routing info in slave BNG:

```
A:Dut-C>config>redundancy# show srrp 1
===============================================================================
SRRP Instance 1
===============================================================================
Description         : (Not Specified)
Admin State         : Up                Oper State       : initialize
Preempt             : yes               One GARP per SAP : no
Monitor Oper Group  : None
System IP           : 10.20.1.3
Service ID          : VPRN 3
Group If            : grpif             MAC Address      : 1c:87:ff:00:00:00
Grp If Description  : N/A
Grp If Admin State  : Up                Grp If Oper State: Down
Subscriber If       : subif
Sub If Admin State  : Up                Sub If Oper State: Up
Address             : 10.11.1.3/16      Gateway IP       : 10.11.1.1
Redundant If        : redundant-interfa*
Red If Admin State  : Up                Red If Oper State: Up
Address             : 10.10.30.3/24
Red Spoke-sdp       : 32:1000
Msg Path SAP        : pw-2:1000
Admin Gateway MAC   :                   Oper Gateway MAC : 00:00:5e:00:01:01
Config Priority     : 1                 In-use Priority  : 1
Master Priority     : 1
Keep-alive Interval : 1 deci-seconds    Master Since     : 05/29/2012 07:22:26
```

```
Master Down Interval: 0.000 sec (Expires in 0.000 sec)
Fib Population Mode : all
VRRP Policy 1       : None               VRRP Policy 2    : None
===============================================================================
* indicates that the corresponding row element may have been truncated.

A:Dut-C>config>redundancy# show service id 3 arp-host
===============================================================================
ARP host table, service 3
===============================================================================
IP Address      Mac Address       Sap Id              Remaining           MC
                                                      Time                Stdby
-------------------------------------------------------------------------------
10.11.1.11      00:80:00:00:00:01 [pw-2:11]           03h38m01s           Yes
10.11.1.12      00:80:00:00:00:02 [pw-2:12]           03h38m02s           Yes
-------------------------------------------------------------------------------
Number of ARP hosts : 2
===============================================================================
A:Dut-C>config>redundancy# show router 3 route-table 10.11.1.11
===============================================================================
Route Table (Service: 3)
===============================================================================
Dest Prefix[Flags]                              Type    Proto    Age          Pref
      Next Hop[Interface Name]                                     Metric
-------------------------------------------------------------------------------
10.11.1.11/32                                   Remote  Sub Mgmt 00h22m03s   0
       [redundant-interface]                                          0
-------------------------------------------------------------------------------
No. of Routes: 1
Flags: L = LFA nexthop available    B = BGP backup route available
       n = Number of times nexthop is repeated
===============================================================================
```

9.10.2.2 ESM on PXC-Based PW-Ports

This section provides an example on how to configure PW-Port based capture SAP that is used in ESM. For more information on PXC Based PW-Ports, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Services Overview Guide*.

PXC Configuration

The following is a PXC configuration example:

```
configure
   port-xc
      pxc 1 create
         port 1/1/1
         no shutdown
      pxc 2 create
         port 2/1/1
         no shutdown
```

With this configuration, ports 1/1/1 and 2/1/1 are auto-provisioned in hybrid mode operating as individual loopback ports. The SR OS system will automatically create a pair of sub-ports per PXC. Those sub-ports are by default in shutdown state:

```
configure
   port pxc-1.a
      shutdown
   port pxc-1.b
      shutdown

   port pxc-2.a
      shutdown
   port pxc-2.a
      shutdown
```

For redundancy purposes and/or increased bandwidth, the PXC sub-ports are aggregated in a LAG:

```
configure
    lag 1 create
        port pxc-1.a
        port pxc-2.a
    lag 2 create
        port pxc-1.b
        port pxc-2.b
```

FPE Configuration

A PW-port is associated with PXC via an FPE configuration. FPE configuration facilitates creation of an internal tunnel over PXC. This tunnel is used to map the external PW to the PW-port. For more information on FPE, refer to *7450 ESS, 7750 SR, 7950 XRS, and VSR Interface Configuration Guide*.

The following is an FPE configuration example:

```
configure
    fwd-path-ext
        fpe 1 create
            path xc-a lag-1 xc-b lag-2
            pw-port
```

The association between xc-a/b (cross-connects) and LAG IDs is performed arbitrarily by the operator. For example, it is allowed to associate xc-a with lag-id 2 (which includes PXC sub-ports on the .b side) and xc-b with lag-id 1 (which includes PXC sub-ports on the .a side). However, the FPE will always assign the .a side of the pxc sub-ports to the transit side of the cross-connect, while the .b side of the pxc subport will be assigned to the termination side of the cross-connect. Refer to the PXC Based PW-Ports section in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN*, for further information about transit/termination side of the cross-connect.

PW-Port Creation

The PW-port must be explicitly created in the SR OS, before mapping between PW and PW-port can be performed.

```
pw-port 100 create
    encap-type qinq
```

SDP Creation for the External PW

The following displays an SDP configuration for the external PW.

```
configure
    service
        sdp 1 create
            signaling tldp
            far-end 10.1.1.1
```

The PW can be static or dynamically signaled, with MPLS or GRE transport.

Mapping Between the External PW and the PW-Port

The stitching of the external PW and PW-port is configured through an Epipe in vc-switching mode.

```
configure
    service
        epipe 10 customer 1 vc-switching create
            spoke-sdp 1:100 create
            pw-port 100 fpe 1
```

Capture PW-SAP Creation

PW-SAPs can be configured as capture SAPs. In this example, a capture PW-SAP with s-tag 3 is created on the pw-port 100.

```
configure
    service vpls 2 create
        sap pw-port-100:3.* capture-sap create
```

From here, ESM functionality is applied to the PW-SAP in the same manner as on any other regular SAP.

Cross-Connecting SAPs to PW-Ports

In addition to PW termination, a PW-port can become a terminating point for a regular SAP. For example:

```
configure
    service
```

```
epipe 10 customer 1 vc-switching create
    sap 1/1/1:10.* create
    pw-port 100 fpe 1
```

In this example, the outer VLAN tag 10 in the payload is removed on ingress and the payload is delivered to the PW-Port where it can be mapped to a capture PW-SAP. This scenario allows traffic distribution from a single I/O port to different EMS termination points (anchor line cards) based on outer VLANs.

9.11 Logical Link Identifier (LLID)

This feature enables service providers to track subscribers on the basis of a virtual-port known as logical line ID (LLID). The LLID (an alphanumeric string) is a logical identification of a subscriber line. Mapping of physical line of a subscriber to LLID is performed via pre-authentication with a separate AAA server than the AAA server used for authenticating the subscriber session during normal access authentication.

LLID serves the purpose of abstracting the physical line of the user from the ISP. If the user moves to a new physical line, the RADIUS server database maintaining the physical line of the subscriber to LLID is updated. Because a subscriber's LLID remains same regardless of subscriber's physical location, using LLID gives service provider a stable and secure identifier for tracking subscriber.

The local user database assigned to the PPPoE node under the group interface can have both a pre-authentication policy and an authentication policy. The purpose of the pre-authentication policy is to retrieve the LLID from the AAA server. The pre-authentication will only extract the calling-station-id attribute (0x31) which is used as the LLID, anything else returned during pre-authentication are simply ignored. If the pre-authentication is missing the LLID, the session will move on to the authentication policy. In the authentication policy that follows, it is possible to use the LLID as the calling-station ID.

It is possible to convey LLID from the LAC to the LNS. The LLID is retrieved through PPPoE pre-authentication where the returned RADIUS attribute calling station ID is used as the LLID. This LLID is selectable attribute in L2TP as a calling-number (AVP 22) to be passed from LAC to LNS. At the LNS, the subscriber calling station number is retrieved from AVP 22 and can be included as an attribute during authentication.

9.12 PADI Authentication Policy for Managed SAP (MSAP)

The PADI Authentication Policy feature enables PADI authentication that retrieves MSAP parameters before pre-authentication and PPPoE authentication.

With this feature, authentication occurs in the following manner.

1. PADI authentication and MSAP authorization:
    1. A capture SAP receives a PPPoE PADI packet.
    2. A check verifies that the LUDB is configured as **pppoe-user-db** on the capture SAP.

       The LUDB must be configured on the capture SAP. Without the LUDB specified, the existing functionality is performed for MSAP creation.
    3. The LUDB host entry is matched to the entry that has the PADI authentication policy.
    4. RADIUS authentication and MSAP authorization occurs.
    5. An MSAP is created if the authentication policy is not configured for PAP/CHAP.
2. (Optional) LLID pre-authentication:
    1. Look up the LLID pre-authentication policy under the LUDB host entry.
    2. Perform RADIUS pre-authentication to obtain an LLID.
3. LCP authentication and L2TP tunnel authorization:
    1. Trigger RADIUS authentication defined by the authentication policy configured in the LUDB host entry in the previous step.
    2. Create the MSAP, finish the LCP authentication with the PPP client, and establish an L2TP tunnel/session.

Figure 148 **Triple Authentication with MSAP Authentication Policy**

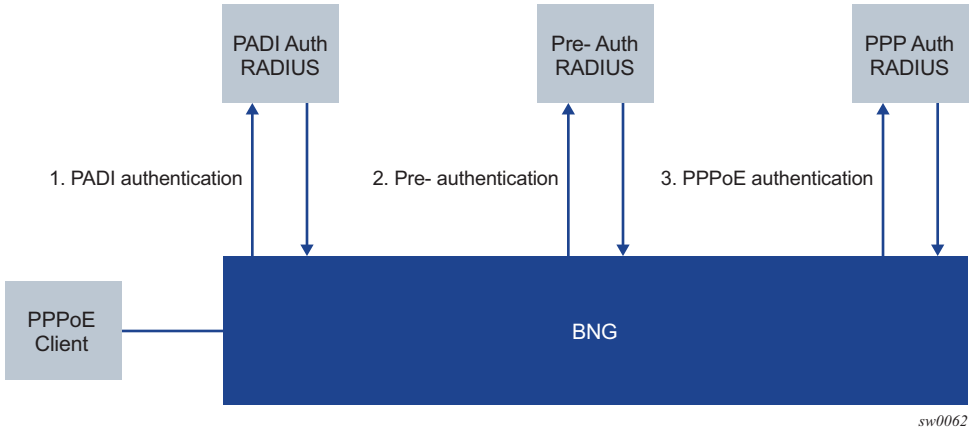

The CLI command **config>subscr-mgmt>loc-user-db>ppp>host>padi-auth-policy** configures a PADI authentication policy used for PADI authentication with MSAP authorization as shown in the CLI example below:

```
configure
    subscriber-mgmt
        local-user-db "ludb-1" create
            ppp
                host "default" create
                    padi-auth-policy "padi_auth"
```

9.13 Open Authentication Model for DHCP and PPPoE Hosts

9.13.1 Terminology

- LUDB – Local User Database configured within the 7450 ESS and 7750 SR.
- IP Address Assignment via DHCP Relay — IP address assignment request (DHCP or IPCP) from the host is relayed to an internal or external DHCP server. Gi-address must be present in this relayed request while the pool name is optional. The internal DHCP server may select the IP address from is local pool based on the gi-address or based on the pool-name present in the request. The IP address selection method is configuration dependent. Third party DHCP servers may consider additional fields in IP address selection process (mac address, circuit-id, and so on).
- IP Address Assignment via DHCP Proxy — A preconfigured IP address in LUDB or RADIUS server is handed out to the host via a DHCP proxy function. This proxy function responds natively using DHCP protocol to the IPoE host. Although PPPoE hosts are not utilizing DHCP protocol, the DHCP proxy functionality within the server is still needed for successful IP address delegation to PPPoE hosts.

9.13.2 LUDB and RADIUS Access Models

During the subscriber-host instantiation phase, various parameters for the hosts are gathered from a single or multiple sources. These parameters represent the level of service within the router to which the host is entitled. Some of the parameters are mandatory for subscriber instantiation while others are optional. The following lists the parameter sources in the order of priority:

- LUDB
- RADIUS
- DHCP Server — The DHCP server directly queries LUDB
- DHCP option processed on DHCP ACK that is indicated in subscriber identification policy.
- Extraction from the DHCP Ack via Python (IPv4 only)
- Defaults that are statically configured on the node (SAP, msap-policy, capture-sap, and subscriber-identification-policy).

In most cases, the host IP address assignment process is controlled by the parameters returned via LUDB or RADIUS. As such, the IP address delegation is integral part of the host instantiation process and will consequently be described in the following sections.

9.13.3 No Authentication

IPoE and PPPoE v4/v6 hosts on static SAPs can be instantiated without the need to access LUDB or RADIUS server. In this case, the default subscriber host parameters (sla-profile, sub-profile, subscriber-id) must be provisioned statically under the SAP. The IP address assignment is provided by internal or external DHCP server. The IP address selection on the router based DHCP server is based on the gi-address while third party DHCP servers may provide additional means to select the IP address (*mac-address*, *circuit-id*, and so on).

A DHCP pool name cannot be provided by an SR-series router DHCP relay agent, since the LUDB and/or RADIUS are not utilized.

This model does not support IP address delegation via DHCP Proxy function since there is no LUDB or RADIUS server available that can supply pre-configured IP address.

Host instantiation without LUDB or RADIUS access on dynamic VLANs (capture SAP and consequently mSAP) is not supported.

9.13.4 LUDB Only Access

Subscriber-host authentication, identification and IP address assignment can be performed via LUDB without the need to access the RADIUS server.

The LUDB is normally configured under the **group-interface>ppp**/**dhcp** hierarchy and can provide subscriber-identification parameters as well as IP addressing parameters:

Pool names for DHCP relay function (IPv4, IPv6 IA-NA, IPv6 IA-PD)

Fixed IP addresses – IPv4, IPv6 IA-NA, IPv6 IA-PD and IPv6 SLAAC prefix.

In case of capture SAP, the LUDB name configured under the capture SAP must much the LUDB name under the **group-interface>dhcp**/**ppp** hierarchy. If the LUDB names do not match, the subscriber-host instantiation will fail.

9.13.5 LUDB Access via DHCPv4 Server

In case that the IPv4 addressing assignment is facilitated by the DHCPv4 relay and an internal DHCPv4 server, the DHCPv4 server itself can query the LUDB for IPv4 address information. LUDB can provide a v4 pool name and IPv4 DHCP options to the DHCPv4 server or it can instruct it to use the gi-address as the IPv4 address selection mechanism.

ESM strings can also be provided via LUDB queried by the DHCPv4 server.

If LUDB access via DHCPv4 server is provided in addition to other authentication means (another LUDB under the group-interface, or RADIUS server), the ESM strings from the LUDB under the grp-interface or from the RADIUS server will have priority over the ESM strings configured under the LUDB accessed by the DHCPv4 server. On the other hand, the IPv4 addressing information will have the highest priority from the LUDB accessed directly by the DHCPv4 server.

Accessing LUDB directly via DHCPv4 server should be used in rare and exceptional cases.

LUDB access under the group-interface, possibly complemented by the RADIUS server will provide necessary means for subscriber-host instantiation in majority of use cases.

9.13.6 RADIUS Only Access

Similar to LUDB-only access, RADIUS server can provide all the necessary information for subscriber-host instantiation, including the IP addressing parameters (pool names or IP addresses/prefixes). Authentication-policy which defines the RADIUS access must be applied to the group-interface.

In case of capture SAP, the authentication policy must be applied under the capture SAP. This authentication policy name must match the authentication policy name that is configured under the group-interface. Otherwise, the host instantiation will fail.

9.13.7 Consecutive Access to LUDB and RADIUS

LUDB and RADIUS access can be combined during subscriber-host instantiation phase.

Configuration wise, LUDB must be referenced under the **group-interface>dhcp/ppp/pppoe** hierarchy (and possibly under the capture SAP), while the authentication-policy is specified within the LUDB. In this fashion, LUDB access is followed by RADIUS access. The subscriber-host parameters retrieved from both sources are combined with LUDB parameters being prioritized over RADIUS parameters in case that both sources return the same parameters.

In case that LUDB and authentication policy are configured simultaneously under the group-interface (and possibly under the capture SAP), the RADIUS authentication policy will be evaluated and LUDB will be ignored.

9.13.8 RADIUS Fallback

If RADIUS server is not accessible (non-responsive), the host instantiation phase can be:

- Terminated if there is no fallback action within authentication policy specified.
- Continued within LUDB if the fallback action within the authentication-policy references LUDB.
- Continued without any response from RADIUS. Subscriber-host will be instantiated if defaults parameters are statically configured or the instantiation will fail in case that the defaults are not available.

The fallback action takes effect once the preconfigured RADIUS timeout period expires.

RADIUS fallback is not supported for DHCPv6 hosts for non-IPoE sessions but is supported for IPoE sessions.

9.14 Flexible Subscriber-Interface Addressing (Unnumbered Subscriber-Interfaces)

9.14.1 Terminology

Subscriber host — A representation of an external host requesting a service. Each such host is fully instantiated within the 7450 ESS and 7750 SR for the purpose of providing traffic control and billing services (for example, QoS, filtering, antispoofing, accounting). The external hosts may represent variety of devices such as regular PCs, STBs, residential gateways, CPEs, VoIP devices. In most cases, the external host will run a DHCPv4/v6 or PPPoEv4/v6 client. DHCP and PPPoE initiation messages from such clients will trigger host instantiation within the router. For this the subscriber host term can be interchangeably used with a term DHCP client or PPPoE client.

9.14.2 Flexible Subscriber-Interface Addressing for IPOE/ PPPoE v4/v6 Subscribers

In certain wholesale or retail environments, the wholesale provider that own the 7450 ESS and 7750 SR BNG does not know the IP addresses that the retailers will assign to their clients in advance. For this reason, wholesaler’s BNG must accept any IP address from retailers and consequently pass it to the client during subscriber-host initiation phase.

Figure 149 **Use Case for Flexible IP Addressing Model**

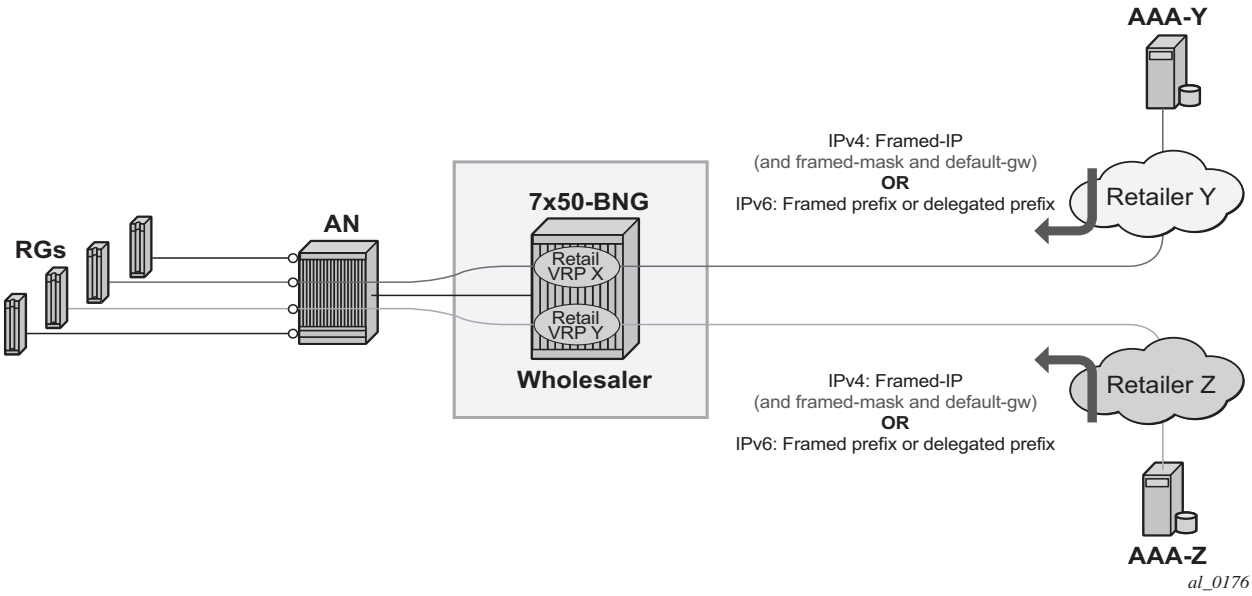

Flexible addressing of the subscriber-interface assumes two deployment scenarios:

1. Subscriber-interface is unnumbered — For example, there is no explicit assigned IP address. Instead the subscriber-interface borrows the IP address from an existing interface that is operationally UP and is located in the same routing instance (router | vprn).

   Note that an interface must have an IP address assigned in order to be operationally UP. Therefore, an unnumbered subscriber-interface must reference another existing interface that is operationally UP in the same routing instance. The subscriber-interface will borrow the IP address from the referenced interface.

   In this case any IP address can be assigned to the subscriber host under the unnumbered subscriber-interface. The subscriber IPv4 address will be installed in the FDB as /32 route while IPv6 address will be installed as an entry of the length anywhere between 64 and 128 bits.

2. Subscriber-interface is numbered — The IP address/prefix is explicitly configured and solely owned by the subscriber-interface.

   In this case, all subscriber IP addresses/prefixes that fall under the subnet/prefix dictated by the configured subscriber-interface IP address/prefix will be directly aggregated under the subscriber-interface subnet. As such they will occupy a single entry in the FDB. The rest of the subscriber hosts with IP addresses/ prefixes that fall outside of the configured range will be installed in the FDB as individual entries (/32 for IPv4 and an entry of the length anywhere between 64 and 128 bits for IPv6 hosts).

9.14.3 Default Gateway in IPv4 Flexible Addressing

In scenarios where subscriber host IPv4 address lies within the configured subscriber-interface subnet, the default-gw IPv4 address for the host will be one of the subscriber-interface IPv4 addresses. In this case, the service provider is aware of the IPv4 addressing scheme in the BNG and as such it will supply the DHCP client with the appropriate default-gw IPv4 address via LUDB, RADIUS or DHCP Server (in that order of priority).

In scenarios where the retail service provider wants to maintain independence from the IPv4 addressing scheme deployed in the BNG (that is controlled by wholesaler), the retailer can always supply its own IPv4 address, the subnet mask and the default-gw IPv4 address. But if the default-gw IPv4 address and/or subnet mask is not supplied by the retailer, then they will be auto-generated by the BNG. Once the default-gw IPv4 address is auto-generated, it will be sent to the requesting DHCP client via DHCP offer in option 3 (RFC 2132, Router Option, section 3.5). There is no additional configuration needed for this action. The BNG will automatically detect whether the default-gw IPv4 address is supplied via LUDB, RADIUS or DHCP server and it will act correspondingly.

The default-gw IPv4 address is auto-generated based on the assigned IPv4 address/mask by setting the last bit of the assigned host IPv4 address to binary 01 or binary 10. For example if the subscriber host's assigned IPv4 address is 10.10.10.10 255.255.255.0, then the default-gw IPv4 address is set to 10.10.10.1. If the assigned IPv4 address is 10.10.10.1 255.255.255.0, then the auto-generated default gateway IPv4 is set to 10.10.10.2.

The default gateway IPv4 address will always have to be within the subscriber's subnet. If it is not, the behavior might be inconsistent. For example:

1. RADIUS (or DHCP) returns IP@, mask and def-gw:
   - IP 10.10.10.1
   - Subnet mask 255.255.255.0
   - Def-gw 10.10.0.254

The subscriber will be successfully instantiated in the BNG but the client may not ARP for a default-gw outside of its configured subnet. Whether the client will or will not ARP for a default-gw outside of its configured subnet will depend on the implementation in the RG and CPE.

- RADIUS returns IP@ and subnet mask.
  - In this case the auto-generated default-gw IPv4 address will always be within subscriber's subnet.

Flexible IPv4 addressing with auto-generated default-gw is supported only in Routed Central Office (RCO) model with routed residential gateways (RGs) or CPEs. In RCO model with bridged residential gateways or CPEs, the default-gw IPv4 addresses and the assigned IPv4 addresses may overlap. Once the IPv4 address of the default-gw is auto-generated, it is possible that the second host behind the bridged residential gateway or CPE is assigned the same IPv4 address as the IPv4 address of the default gateway of the first host. Such hosts would not be able to communicate with outside world.

For example:

RADIUS or DHCP server assigns IPv4 address and subnet mask to the first host in a bridged environment:

IP1: 10.10.10.1

Auto-generated default-gw IPv4 address: 10.10.10.2

Since the RADIUS and DHCP Server are not aware of the auto-generated default-gw, they may assign the following IPv4 address to the second host that comes on-line:

IP 2: 10.10.10.2 (same IPv4 address as the default-gw IPv4 address of the first host)

Auto-generated default-gw IPv4 address: 10.10.10.1

Now the first host will forward all traffic outside of the configured subnet to the second hosts which will discard this traffic, effectively rendering this operation model non-deployable. And vice versa.

9.14.4 IPv4 Subnet Sharing

Subnet sharing between the hosts in flexible IPv4 addressing model is supported. In other words, in flexible IPv4 addressing model the operator can assign all IPv4 addresses (minus one – the default-we IPv4 address) from a given subnet. In this fashion, all subscribers (routed RGs or CPEs) within a single subnet can share the same default gateway.

For example if the operator owns the IPv4 subnet 10.10.10.0/24, then one IPv4 address can be set aside for the default-gw (for example 10.10.10.254) and the remaining addresses can be assigned to the subscriber (routed RGs or CPEs). An example would be:

RG1: IP=10.10.10.1/24 def-gw 10.10.10.254

RG2: IP=10.10.10.2/24 def-gw 10.10.10.254

RG3: IP=10.10.10.3/24 def-gw 10.10.10.254

:

RG100: IP=10.10.10.100/24 def-gw 10.10.10.254

The subnet sharing is also supported in conjunction with auto-generated default-gw IPv4 address. The implication of this is that the IPv4 address of the default-gw can collide with the same IPv4 address already assigned to an existing subscriber. This is not an issue for routed RGs or CPEs since the BNG will always answers ARPs for the IPv4 address of the default-gw with its own MAC address. However, local-proxy ARP functionality in the 7450 ESS and 7750 SR BNG must be enabled to support this.

This behavior can be further clarified with the following example.

Let’s assume that we have scenario with two routed RGs:

RG-1, IP=10.10.10.0/24, default-gw IP=10.10.10.1

RG-2, IP=10.10.10.1/24, default-gw IP=10.10.10.0

Once RG-1 ARPs for its default gateway of 10.10.10.1, the BNG will reply with its own MAC address.

Now that host RG-1 has resolved ARP for it default-gw (MAC address pointing to the router), it can send traffic to the outside world via the BNG. When such traffic arrives to the router, the destination IPv4 address of the received packet will determine the forwarding decision within the router. If the destination IPv4 address matches the IPv4 address of any subscriber (RG) instantiated within the system, the traffic will be forwarded to the that RG. This also includes the case where the destination IPv4 address is the default-gw IPv4 address (10.10.10.1), which represents just another RG within the router. The traffic will be consequently passed from RG-1 via 7450 ESS and 7750 SR to RG-2.

9.14.5 IPv4 Subnet Mask Auto-Generation

The subnet mask corresponding to the IPv4 address assigned to the subscriber is auto-generated in case that the IPv4 addressing authority (LUDB, RADIUS or DHCP Server) does not supply it. The subnet mask is derived from the IPv4 address of the subscriber and possibly the default-gw IPv4 address and it is the smallest subnet that contains both, the IPv4 address of the subscriber and the default-gw.

For example if the RADIUS received IPv4 address is 10.10.10.138 and the received default –gw IPv4 address is 10.10.10.170, then the subnet mask will be auto-generated and set to 255.255.255.192 (/26).

138 = 10001010

170 = 10101010

192 = 11000000

In case that neither the subnet mask nor the default-gw are returned, then both would be auto-generated:

1. Subnet mask would be set to /31
2. Default-gateway which must belong to the subscriber's subnet would be set to 10.10.10.139.

In cases where the host IPv4 address and the default-gw are directly supplied by the addressing authority but the subnet mask is missing, the subnet mask auto-generation may cause the host part of the default-gw IPv4 address to become a broadcast IPv4 address. If this is an issue, then it can be avoided by directly providing the subnet mask via the addressing authority.

9.14.6 Local-proxy-arp and arp-populate

Local-proxy-arp and arp-populate are two commands that are relevant only to IPoEv4 hosts.

The **local-proxy-arp** command ensures that the router answers ARP Requests with its own MAC address for any active IPv4 address under the subnet on which the ARP request arrived. The active IPv4 address is considered the one that is assigned to an already instantiated hosts or the default-gw (even auto-generated).

In absence of **local-proxy-arp** command, the only ARP Request that the router will answer is the one for the statically configured IPv4 addresses of the subscriber-interface. In flexible IPv4 addressing, the IPv4 address of the default-gw does not necessarily match any of the configured subscriber-interface IPv4 addresses. The ARP Request for such default-gw IPv4 address would go unanswered. Consequently, the subscriber hosts would not be able to communicate with outside world. Therefore, the flexible IPv4 addressing requires that the **local-proxy-arp** command is configured.

The **arp-populate** command disables dynamic learning of ARP entries (IPv4<->MAC mapping) on an interface based on the ARP protocol. In this case, the ARP table is populated based on the DHCPv4 lease state table which contains IPv4<->MAC mappings obtained via DHCP processing during the host instantiation phase. Arp-populate functionality is highly desirable in case of flexible IPv4 addressing.

When the **arp-populate** command is disabled the ARP entries are dynamically learned based on the ARP protocol. This, in conjunction with flexible IPv4 addressing may cause certain issues. Consider the following example:

- The subscriber host is instantiated in the 7450 ESS and 7750 SR
- The subscriber interface is unnumbered
- The ARP table does not contain an ARP entry for the subscriber-host

In this case, downstream traffic towards the subscriber host will trigger the router to send ARP Request for the subscriber host IPv4 address. The router needs to know the MAC address of the subscriber-host in order to forward traffic. Since the subscriber-interface is unnumbered, the source IPv4 address of the ARP request is unknown and consequently the ARP request will not be sent. As a result, downstream traffic is dropped.

Note however that the above example is an unlikely scenario. If the subscriber host sends the ARP request for the default-gw first, the router would create an entry in the ARP table for it and the issue would be resolved. This is the most likely outcome since the subscriber host will always try to initiate communication with the outside and therefor ARP for the IPv4 address of the default-gw (which is a 7450 ESS and 7750 SR).

9.14.7 Gi-address Configuration Consideration

With flexible IPv4 address assignment, the gi-address can be configured as any IPv4 address that is already assigned to an interface (loopback interface, regular interface attached to physical port or subscriber interface) within the same routing instance (VRF or GRT).

9.14.8 PPPoE Considerations

PPPoE subscriber hosts do not have the concept of default-gw. Consequently the default-gw auto-generation concept does not apply to PPPoE hosts.

9.14.9 IPoEv6 Considerations

The default-gw for IPoEv6 hosts is link-local IPv6 address. Since this address is always present, there is no need for auto-generation during the subscriber instantiation time.

SLAAC hosts are installed as /64 entries, the length of the installed DHCP-PD prefix is dictated by the prefix-length and the DHCP-NA hosts are installed as /128 entries.

9.14.10 General Configuration Guidelines for Flexible IP Address Assignment

Flexible IP addressing for IPoE/PPPoE v4 and v6 hosts is by default disabled. In other words, the subscriber hosts will be instantiated in the BNG with ability to forward traffic only if their assigned IP addresses belong to one of the configured subnets/prefixes that are associated with the subscriber-interfaces. IPv4 and IPv6 cases will be examined separately:

IPv4:

By default, IPoE and PPPoE subscriber host creation fails in the following two cases:

1. The subscriber-interface does not have an IPv4 address configured, and therefore it is be operationally down. This configuration is also known as unnumbered subscriber-interface.
2. The subscriber-interface does have an IPv4 address configured but the IPv4 address assigned to the subscriber host itself is outside of the subscriber-interface configured subnet(s). In such case the host will be instantiated but the forwarding is disabled.

Subscriber host instantiation and forwarding can be explicitly enabled for both cases above with flexible IP addressing functionality.

For case 1, this can be achieved by borrowing an IP address for the subscriber-interface from any interface that is operationally up within the given routing context. This functionality can be enabled with the following command:

```
configure service ies <id>
configure service vprn <id>
subscriber-interface <intf-name>
unnumbered <ip-addr | interface-name >
```

To enable forwarding for the subscribers whose IP address falls outside of the configured subnet under the subscriber-interface (case 2), the following command must be entered:

```
configure service ies <id>
configure service vprn <id>
    subscriber-interface <intf-name>
        allow-unmatching-subnets
```

The above commands (**unnumbered** and **allow-unmatching-subnets**) are mutually exclusive. In addition, the **unnumbered** command can be configured only if the subscriber interface does not have an IP address already configured. Otherwise the execution of this command will fail.

In both of these cases the host will be installed in the routing table as /32.

IPv6:

For IPv6 there is a single command that will enable flexible IP addressing for both cases:

1. IPv6 prefixes are not configured under the **subscriber-interface>ipv6** node
2. IPv6 prefixes are configured but the actual address or prefix assigned to the subscriber (via DHCP, LUDB or RADIUS) is outside any prefix that is configured under the **subscriber-interface>ipv6** hierarchy.

This single command is:

```
configure service ies <id>
configure service vprn <id>
    subscriber-interface <name>
        ipv6
        allow-unmatching-prefixes
```

To summarize, the following scenarios are possible:

- PPPoEv4
  - An IPv4 address under the subscriber-interface is configured
    - By default hosts outside of the sub-intf subnet are instantiated but they are in a non-forwarding-state. Traffic is dropped.
    - **allow-unmatching-subnets** is configured. This command is allowed only if subscriber-interface has also configured its own IPv4 address(es). In this case the IP address for IPCP negotiation is one of the sub-intf addresses. Hosts outside of the sub-intf subnets are instantiated and forwarded.

- The **unnumbered** <*ip-address* | *intf*> command is not allowed in this scenario.

– An IPv4 address under the subscriber-interface is not configured

- By default, the subscriber-interface is operationally down. Subscribers cannot be instantiated.
- The **allow-unmatching-subnets** command has no effect since subscriber-interface does not have an IPv4 address configured and is therefore operationally down. No subscribers can be instantiated.
- The **unnumbered** <*ip-address* | *intf*> command is the only viable option in this case. The subscriber-interface borrows an IPv4 address from another interface that is operationally UP and consequently this allows subscribers to be instantiated. This command is mutually exclusive with **allow-unmatching-subnets**. In addition, this command can only be configured if the subscriber-interface itself does not have explicitly configured an IPv4 address.

- IPoEv4

Similar to the PPPoE case above.

- IPoEv6 and PPPoEv6 — the **allow-unmatching-prefixes** command is independent of any IPv4 command related to flexible IP address assignment (**unnumbered** or **allow-unmatching-subnets**). This command can always be enabled, regardless of the v6 prefixes configured under the **subscriber-interface>ipv6** hierarchy. Any subscriber, regardless of the subscriber-interface prefix configuration will be instantiated and forwarded.

9.14.11 Caveats

- Auto-generation of the default-gw IPv4 address is supported only in RCO model with routed RGs/CPEs. Bridged RGs/CPEs are not supported.
- A configured IPv4 address cannot be removed from the subscriber-interface when DHCPv4 hosts under the corresponding subnet are instantiated in the system.
- An IPv4 address cannot be configured under the subscriber-interface while (unnumbered) DHCPv4 hosts under that subnet are already instantiated.
- Executing the **no allow-unmatching-subnets** command is only allowed when there are no unnumbered DHCPv4 hosts instantiated under the subscriber-interface.

9.15 uRPF for Subscriber Management

uRPF is supported for IPv4 and IPv6 dual-stack subscribers with framed routes.

For IPv4, uRPF is supported on group interfaces using anti-spoofing filters. A group interface configured for NATed subscribers will be configured with MAC/IP/PPPoE Session-ID anti-spoofing filters.

IPv6 subscribers, which are non-NAT, are always treated as being on a local subnet. For such subscribers, a BNG will install a FDB entry for local routes that match either the wan-host prefix, or the delegated prefix, or both. In strict mode for IPv6 ESM, the uRPF check will check not just that the route matching the SA (which should be a local route, i.e. a subnet) would route the packet back out of the interface it came in on, but in addition that we would route the packet out to the same SAP it was received on.

SR OS supports the ability to configure a NH-MAC anti-spoof type for non-NATed subscribers. When configured, the datapath performs ingress anti-spoofing based on source MAC address and egress anti-spoof (also referred to as egress subscriber-host look-up) based on the nh-ip address.

The NH-MAC anti-spoof type is configured under the following context:

config>service>vprn>if>sap

config>service>ies>sub-if>grp-if>sap

config>service>vprn>sub-if>grp-if>sap

config>subscr-mgmt>msap-policy

A uRPF check is also performed that prefixes delegated to a subscriber on that MAC address exist in the FDB.

9.16 IPoE Sessions

The IP stacks of dual stack IPoE devices are set up and configured independently using different protocols such as DHCPv4, DHCPv6 or SLAAC. As opposed to PPPoE, there is no single protocol that binds the IP stacks from a single end device together.

To facilitate subscriber management of dual stack IPoE devices as a single entity similar as for PPPoE sessions instead of handling individual IPoE subscriber hosts, there is a need for a logical IPoE session construct. An IPoE session enables single authentication, session accounting and policy management (mid-session changes) for dual stack IPoE devices.

An IPoE session is a logical grouping of IPoEv4 and IPoEv6 subscriber hosts that represent the different IP stacks of a single end device and that share authentication data such as subscriber id, subscriber and SLA profile, session-timeout, and so on. The grouping of subscriber hosts in an IPoE session is based on a configurable session key per group-interface. The IPoE session key includes by default the SAP identifier and MAC address and can be extended with Circuit-Id/Interface-Id or Remote-Id. For DHCPv6 Remote-Id, the enterprise number is excluded from the session-key. Circuit-id/Interface-Id or Remote-id should only be used in the IPoE session key if all subscriber host associated with the IPoE session have this field in their protocol trigger packets. The IPoE session creation or subscriber host association to an IPoE session fails if the Circuit-Id/Interface-Id or Remote-id is not present in a trigger packet while the field is part of the session-key.

Figure 150 **IPoE Session**

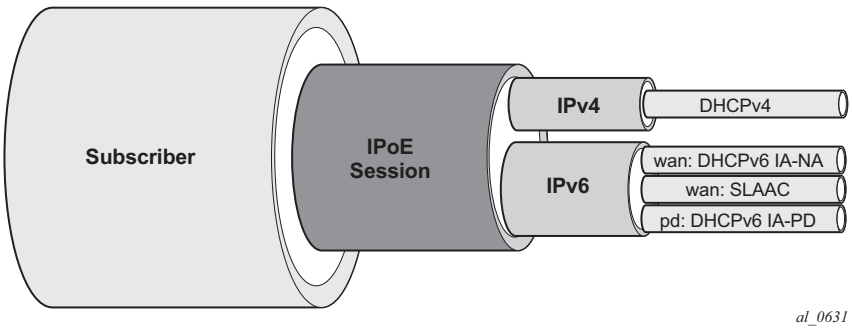

An IPoE session represents a single end device and can have following associated IP stacks:

- IPv4 — A single DHCPv4 host.

- IPv6 WAN — One DHCPv6 IA-NA host and/or one SLAAC host.
- IPv6 PD — One DHCPv6 IA-PD host or PD as managed route.

A violation of the above rules will result in a setup failure of the subscriber host when an attempt is made to associate it to the IPoE session.

9.16.1 Enabling IPoE Sessions

IPoE sessions are supported in a Routed CO environment with Enhanced Subscriber Management (ESM) enabled. To enable the IPoE session instantiation, the **ipoe-session** CLI context on the capture SAP (managed SAP scenario) and/or group-interface must be configured to no shutdown. See also the configuration steps below.

Important Notes:

- Enabling IPoE sessions on a group-interface with active subscriber hosts will trigger a migration. Use one of the following CLI command to determine if there are active hosts on a group-interface:
  - Check the number of subscriber hosts on a group-interface:

    Note that DHCPv6 IA-PD modeled as a managed route are not counted with this command.
    show service id <service-id> subscriber-hosts detail | match <group-int-name> | count
  - Check the number of IP stacks (Client types) attached on a group-interface.
    tools dump router <router-instance> ipoe-session migration interface <group-int-name>

    The active number of IP stacks (Client types) on the group-interface are listed per type as well as if they are associated with an IPoE session or not.

    Note that DHCPv6 IA-PD modeled as a managed route is also counted in the DHCPv6 type counter.
  - Check the number of DHCPv4 leases, DHCPv6 leases and SLAAC hosts on the group-interface that are not attached to a session:
    # show service id <service-id> dhcp4 lease-state interface <group-int-name> session none
    # show service id <service-id> dhcp6 lease-state interface <group-int-name> session none
    # show service id <service-id> slaac host interface <group-int-name> session none

    Note that DHCPv6 IA-PD modeled as a managed route is also counted in the DHCPv6 lease state counter.

If there are active hosts on the group-interface, make sure you have read the “IPoE session migration” section prior to enabling IPoE sessions.

- Disabling IPoE sessions by executing an **ipoe-session shutdown** command or **no ipoe-session** command on a group interface will delete all active sessions and associated hosts on that group interface, resulting in service impact for these subscribers. Use one of the following CLI commands to determine if there are active ipoe-sessions on a group-interface:

```
# show service id <service-id> ipoe session interface <group-int-name>
# tools dump router <router-instance> ipoe-session migration interface <ip-int-name>
```

If there are active IPoE sessions on the group-interface, be aware that disabling IPoE sessions on the group-interface will result in service impact for those sessions.

9.16.2 IPoE Session Authentication

A single authentication is performed for all subscriber hosts that belong to the same IPoE session. The table below lists the packets that trigger an IPoE session authentication.

Table 89 **IPoE Session Authentication Trigger Packets**

IP stack	Trigger packets
IPv4	DHCPv4 Discover
	DHCPv4 Request
IPv6 WAN	DHCPv6 Solicit
	DHCPv6 Request
	DHCPv6 Relay Forward (Solicit)
	DHCPv6 Relay Request (Solicit)
	Router Solicitation
IPv6 PD	DHCPv6 Solicit
	DHCPv6 Request
	DHCPv6 Relay Forward (Solicit)
	DHCPv6 Relay Request (Solicit)

When a trigger packet is received on a capture SAP or group-interface with IPoE sessions enabled, an IPoE session lookup is performed based on the configured IPoE session key:

- If no IPoE session is found, a new session is created and authenticated following the ESM authentication configuration such as local user database lookup, Radius or Diameter authentication, defaults, and such. After successful authentication, the authentication data is stored in the IPoE session state. The subscriber host is created and associated with the session.
- If an IPoE session already exists, and no re-authentication must be performed then the subscriber host is created using the stored IPoE session data. The subscriber host is associated with the session.
- If an IPoE session already exists, and re-authentication must be performed then the session is re-authenticated. When successful, the authentication data for the IPoE session is updated and applied to all associated hosts. The subscriber host is created and associated with the session. When unsuccessful, existing hosts associated with the session are not impacted and the session data is kept unchanged.

Re-authentication is by default disabled for IPoE sessions. To enable re-authentication, a minimum authentication interval must be configured. The min-auth-interval CLI parameter configures the maximum frequency of re-authentications by specifying a minimum interval between two non-forced authentications for the same IPoE session. A re-authentication is triggered by the renewal of any host belonging to the IPoE session. Setting the min-auth-interval to zero seconds, will always re-authenticate on each trigger packet. The **re-authentication** command in a RADIUS authentication policy is ignored for IPoE session authentication.

A forced authentication is performed when the Circuit-Id/Interface-Id or Remote-Id in the trigger packet has changed. An empty or absent Circuit-Id/Interface-Id or Remote-Id is not considered as a change. The default forced authentication behavior is changed with the **force-auth** command in the **group-interface>ipoe-session** context: only force authenticate on Circuit-Id/Interface-Id change or only force authenticate on Remote-Id change or disable forced authentications.

A new local user database config in the ipoe-session CLI context on a capture SAP or group interface ensures that all subscriber hosts associated with an IPoE session are using the same database and therefore common match criteria. The per subscriber host type user-db configurations, such as **ipv6 dhcp6 user-db**, **dhcp user-db** and **rtr-solicit-user-db** are ignored when IPoE sessions are enabled.

9.16.3 IPoE Session Accounting

All Radius accounting modes can be enabled for IPoE sessions: queue instance accounting, host accounting or session accounting.

With session accounting, a Radius accounting start is generated when the first host of the session is created and an accounting stop when the last host of the session is deleted. The generation and interval of periodic interim updates can be configured. Optionally, triggered interim update messages can be generated when a host is deleted from the session or an additional host is associated.

A unique accounting session id is generated for the IPoE session and is used in RADIUS session accounting. The IPoE session accounting session id can be included in the Radius Access Request message via CLI **config>subscr-mgmt>auth-plcy# include-radius-attribute acct-session-id session.**

This accounting session ID can also be used in RADIUS CoA or Disconnect Messages to target the IPoE session.

9.16.4 IPoE Session Mid-Session Changes

Mid-session changes such as those initiated via RADIUS CoA or Diameter Gx RAR are applied to all hosts associated with the IPoE session.

A RADIUS CoA message targeting any host of an IPoE session has the same effect as a Radius CoA message targeting the IPoE session using the IPoE session Acct-Session-Id as key: all host of the session are targeted and the session state is updated with the new data.

The following tools commands are available to manually enforce a mid-session change:

```
# tools perform subscriber-mgmt edit-ipoe-session sap <sap-id> mac <mac-address>
[subscriber <sub-ident-string>] [sub-profile-string <sub-profile-string>] [sla-
profile-string <sla-profile-string>] [inter-dest-id <intermediate-destination-id>]
[ancp-string <ancp-string>] [app-profile-string <app-profile-string>] [circuit-id
<circuit-id>] [remote-id <remote-id>]

# tools perform subscriber-mgmt eval-ipoe-session [svc-id <service-id>] [sap <sap-
id>]
[mac <mac-address>] [circuit-id <circuit-id>] [remote-id <remote-id>] [subscriber
<sub-ident-string>]
```

9.16.5 IPoE Session Termination

When the last subscriber host associated with an IPoE session is deleted from the system, then the IPoE session is also deleted.

An IPoE session and all associated subscriber hosts can be deleted via

- CLI clear command: **clear service id <service-id> ipoe session**
- An **ipoe-session no shutdown** CLI command on a group-interface
- A **no ipoe-session** CLI command on a group-interface. This command resets to the default behavior, which is IPoE sessions disabled.
- Session timeout, configured in the IPoE session policy or obtained from AAA
- Idle timeout
- RADIUS Disconnect Message
- Diameter Gx session termination
- Credit Control: Radius or Diameter Gy

9.16.6 Limiting the Number of IPoE sessions

The number of IPoE sessions per SAP is limited with the **sap-session-limit** command configured in the **group interface ipoe-session** context.

The number of IPoE sessions per group interface or retail subscriber interface is limited with the **session-limit** CLI command configured in the group interface ipoe-session or retail subscriber interface ipoe-session context.

Hosts associated with IPoE sessions are subject to the SLA Profile Instance limits configured in the **sla profile >config>subscr-mgmt>sla-prof>host-limits** context.

9.16.7 SAP Session Index

The system will keep track of the number of IPoE sessions active on a given SAP and assign a per SAP session index to each such that always the lowest free index is assigned to the next active IPoE session. When RADIUS authentication is used, the SAP session index can be sent to, and received from, the RADIUS server using the [26-6527-180] Alc-SAP-Session-Index attribute.

It should only be used in a subscriber per VLAN model as the session index is per SAP.

The SAP session index allows IPoE sessions in a bridged RG environment to have their own set of queues for QoS and accounting purposes when using the same SLA profile name received from a RADIUS server. See Subscriber per PPPoE Session Index for further details.

Alternatively, this can be achieved by configuring per-session SPI sharing in the SLA profile as described in SLA Profile Instance Sharing.

9.16.8 Resiliency

For non-redundant BNG deployments, the IPoE session state is stored in the subscriber-mgmt persistency file for recovery from Compact Flash after a node maintenance operation or failure. This is configured at the system persistence CLI context.

For multi chassis redundancy scenarios, the IPoE session state is synchronized via the “sub-mgmt ipoe” Multi Chassis Synchronization (MCS) application.

9.16.9 Notes

- Static hosts can be configured on a group-interface with IPoE sessions enabled. A static host will not be associated with an IPoE session.
- Up to sixteen Framed-Routes and sixteen Framed-IPv6-Routes can be associated with an IPoE session.
- A fall back action (accept or local user database lookup) when no Radius servers are available for Radius authentication can be specified for IPoE sessions.
- Lawful Intercept sources initiated from Radius always include all IP stacks from the IPoE session regardless the targeted host in the CoA message.
- ARP hosts are not supported in an IPoE session and cannot be instantiated on a group-interface with IPoE sessions enabled.
- The creation of an IPv4 host using the Alc-Create-Host attribute in a Radius CoA message is not supported on a group-interface with IPoE session enabled.
- A local user database host identification based on option60 is ignored when authenticating an IPoE session.
- RADIUS authentication of an IPoE session fails when the **user-name-format** is configured to **mac-giaddr** or **ppp-user-name**.
- The alc.dtc.setESM() API in the DHCP Transaction Cache (DTC) Python module cannot be used in combination with IPoE sessions.

- The DHCP Python module (alc.dhcp) used to derive subscriber host attributes from a DHCPv4 ACK message is not supported in combination with IPoE sessions.
- A RADIUS CoA message containing an Alc-Force-Nak or Alc-Force-Renew attribute is not supported for IPoE sessions
- Subscriber Host Connectivity Verification (SHCV) will continue to work on a per-stack basis. In other words, in a dual stack scenario with SHCV action remove enabled for both stacks, a failure in IPv4 connectivity will not clean up the session unless the IPv4 subscriber host was the last associated host.

9.16.10 Configuration Steps

To create an IPoE session policy:

```
config
    subscr-mgmt
        ipoe-session-policy "ipoe-policy-1" create
            description "Default IPoE session policy"
            session-key sap mac            # default
            no session-timeout             # default
        exit
```

Enable IPoE sessions on the capture SAP and/or group interface.

If IPoE sessions is enabled on a capture-sap, then it must also be enabled on the target group-interface. If an IPoE session local user database lookup is configured at the capture-sap, then the same local user database lookup must be configured at the target group-interface.

```
config
    service

        vpls 10 customer 1 create
            ---snip---
            sap 1/1/4:*.* capture-sap create
                ---snip---
                ipoe-session
                    description "IPoE sessions - capture-sap"
                    ipoe-session-policy "ipoe-policy-1"
                    user-db "ludb-1"
                    no shutdown
                exit

        ies 1000 customer 1 create
            subscriber-interface "sub-int-1" create
                ---snip---
                group-interface "group-int-1-1" create
                    ---snip---
```

```
ipoe-session
    description "IPoE sessions - IES group-interface"
    force-auth cid-change rid-change       # default
    ipoe-session-policy "ipoe-policy-1"
    min-auth-interval infinite             # default
    sap-session-limit 1                    # default
    session-limit 1000
    user-db "ludb-1"
    no shutdown
exit
```

To display the IPoE session state, use following command:

```
# show service id <service-id> ipoe session [detail]
```

9.16.11 IPoE Session Migration

This section is only applicable when enabling IPoE sessions on a group interface with active subscriber hosts. When there are no active subscriber hosts on a group interface, there is no need for a migration. Use one of the following CLI commands to determine if there are active hosts on a group interface:

- Check the number of subscriber hosts on a group interface:
  # show service id <service-id> subscriber-hosts detail | match <group-int-name> | count

  DHCPv6 IA-PDs modeled as a managed route are not counted with this command.
- Check the number of IP stacks (client types) attached on a group interface.
  # tools dump router <router-instance> ipoe-session migration interface <group-int-name>

  The number of active IP stacks (client types) on the group interface are listed per type whether or not they are associated with an IPoE session.

  DHCPv6 IA-PDs modeled as a managed route are also counted in the DHCPv6 type counter.
- Check the number of DHCPv4 leases, DHCPv6 leases, and SLAAC hosts on the group interface that are not attached to a session:
  # show service id <service-id> dhcp4 lease-state interface <group-int-name> session none
  # show service id <service-id> dhcp6 lease-state interface <group-int-name> session none
  # show service id <service-id> slaac host interface <group-int-name> session none

DHCPv6 IA-PDs modeled as a managed route are also counted in the DHCPv6 lease state counter.

By default, IPoE sessions are disabled on a group interface **(ipoe-session shutdown**). Enabling IPoE sessions on a group interface with active subscriber hosts starts a migration process and should be planned carefully to allow a seamless migration.

A migration is required due to the nature of IPoE sessions: a single authentication is performed for all hosts (IP stacks) of a dual-stack end device. All hosts (IP stacks) in an IPoE session share the same MAC address, SAP, and optionally Circuit-ID / Interface-ID or Remote-ID which are configured as the **session-key** in the **ipoe-session-policy**. To determine if hosts (IP stacks) belong to a single session, a new trigger packet is required to obtain the session key.

To guarantee a correct IPoE session configuration and a correct authentication database, the migration is performed when the host state is renewed and a new trigger packet is received:

- DHCPv4 renew or rebind for DHCPv4 hosts
- DHCPv6 renew or rebind for DHCPv6 hosts (IA-NA and IA-PD)
- DHCPv4 renew for IPoE linked SLAAC hosts
- Router Solicit for RS triggered SLAAC hosts

The duration of a migration is therefore dependent on the lease times for DHCPv4 and DHCPv6 hosts and for IPoE linked SLAAC hosts. If possible, the lease times could temporarily be reduced to a couple of hours to facilitate the migration process.

The actual migration is started by the arrival of a new trigger packet of an IP stack (host) that is not associated with an IPoE session. The IPoE session key is composed of the data in the trigger packet (MAC address and SAP, by default). If an IPoE session exists for the obtained IPoE session key, the corresponding session data is used for authentication. If no IPoE session exists for the obtained IPoE session key, authentication is performed, and based on the result, a new IPoE session is created. The old host state is deleted from the system and a trap is sent to indicate that this host is being migrated. A new host (IP stack) is created and associated with the IPoE session. When RADIUS accounting is enabled, this may result in an accounting start and stop depending on the accounting mode. For host accounting, an accounting stop is followed immediately by an accounting start. For queue instance accounting, an accounting stop is generated when the last host associated with the queue instance is migrated. An accounting start is generated when the first host is associated with the IPoE session.

The following notes must be taken into account for the migration procedure:

- For multi-chassis redundant nodes, IPoE sessions should be enabled first on the standby node and immediately thereafter on the active node.
- A renew as part of a DHCPv4 lease split operation will not trigger a migration to the IPoE session. The migration starts only when the renew is forwarded to the DHCP server.
- For DHCPv4 RADIUS proxy scenarios, it is recommended that the lease time be specified the with the [26-6527-174] Alc-Lease-Time RADIUS attribute instead of the [27] Session-Timeout attribute. After migration, the [27] Session-Timeout attribute is interpreted as the number of seconds before the session is terminated.
- DHCPv6 IA-PD modeled as a managed route may migrate separately from the IPv6 SLAAC host it is associated with for its next-hop. This could result in a temporary service impact until both the managed route and next-hop host are migrated.
- The migration of idle Router Solicit SLAAC hosts can be facilitated by specifying an inactivity timer.
- When the subscriber ID is auto-generated (**auto-sub-id**), then a new **sub-id** will be generated after migration. This may result in a temporary increase in used resources such as queues until all hosts from a subscriber are migrated.

Example high-level migration steps.

Important Notes

- It is recommended that a migration plan be built for the target network and validate the plan in advance in a lab environment.
- It is recommended that the migration be performed per group interface or capture SAP with all possible target group interfaces and that the next migration only be started when the previous one is successfully completed.
- When managed SAPs (MSAPs) are used, enabling an IPoE session on a group interface while not enabling IPoE sessions on the corresponding capture SAP, or enabling an IPoE session on a capture SAP while not enabling IPoE sessions on the target group interface, will result in session setup failures for sessions where no MSAPs exist.

1. Using the CLI commands described at the beginning of this section, check if an IPoE session migration is applicable. A migration is not required when there are no active subscriber hosts on the target group interfaces.
2. Check if all preconditions are met:

a. There are no conflicting requirements with IPoE sessions such as ARP host support on the same group interface or local user database authentication based on option 60. Check the Notes section above for a list of possible conflicts.

b. IPoE session configuration is complete on the group interfaces and corresponding **capture-sap**: **ipoe-session-policy** (session-key) and on the optional local user database. On the group interfaces, the IPoE session limits should be configured as needed using the **session-limit** and **sap-session-limit** commands.

c. Authentication servers are up to date to provide all required authentication data for a single dual-stack end device based on single authentication (for example, return both IP address and IPv6 prefix in a proxy scenario).

d. Accounting servers are ready to deal with accounting stop/start when hosts migrate to an IPoE session.

3. Take a snapshot of the active hosts before the migration. Use the commands as described above. The following command provides a summary view: **tools>dump>router <router-instance>>ipoe-session>migration>interface <group-int-name>**

4. Start the migration by enabling an IPoE session on the group interface and for MSAPs, by enabling an IPoE session on the capture SAP.

5. Monitor the progress during migration. Review the events (for example, by using the **show log log-id** 99 command) and check the number of hosts migrated with the CLI command:

   # tools>dump>router <*router-instance*>>ipoe-session>**migration>interface** <*group-int-name*>

   The following event is generated when a host is deleted due to migration:

   4 2015/06/29 19:37:57.47 UTC WARNING: SVCMGR #2559 Base IPoE session "IPoE session migration deleted host 2001:db8:2:101::1 on SAP 1/1/4:1201.2 in service 1000"

   2 2015/06/29 19:37:29.41 UTC WARNING: SVCMGR #2559 Base IPoE session "IPoE session migration deleted host 10.1.1.101 on SAP 1/1/4:1201.2 in service 1000"

   DHCP lease states and SLAAC host states associated with IPoE sessions can be found with:

   # show service id <*service-id*> **dhcp4 lease-state interface** <group-int-name> session ipoe

   # show service id <*service-id*> **dhcp6 lease-state interface** <*group-int-name*> **session ipoe**

   # show service id <*service-id*> **slaac host interface** <*group-int-name*> **session ipoe**

The migration is finished when all hosts are associated with an IPoE session. The counters in the column “Non IPoE session” should be all zero. For example:

```
# tools dump router ipoe-session migration interface group-int-1-1
===============================================================================
 Type session               Total     IPoE session     Non IPoE
===============================================================================
 Group-interface: group-int-1-1 (IPoE session enabled)
-------------------------------------------------------------------------------
 DHCPv4              16384     16384            0
 DHCPv6              16384     16384            0
 SLAAC               4096      4096             0
-------------------------------------------------------------------------------
 IPoE sessions       20480
===============================================================================
```

6. Perform post migration steps. For example, verify that the number of users prior to and after the migration are in the same order of magnitude (users may connect and disconnect during the migration). Enable session accounting if required.

9.16.11.1 Additional Notes for IPoE Session Migration of IPv4 Hosts as a Control Channel for Dynamic Data Services

During the migration of an IPv4 host as a control channel for Dynamic Data Services to an IPoE session as a control channel, the associated dynamic data services will be deleted and recreated based on the IPoE session authentication data.

When IPoE sessions are enabled on the group interface, at the next DHCPv4 renew or rebind:

- the IPv4 host (control channel) is deleted
- the associated dynamic data services are deleted
- the IPv4 host is added in the IPoE session
- new dynamic data services are created based on IPoE session authentication data

9.17 Data-triggered Subscriber Management

This feature allows the creation of ESM subscribers and hosts based on the receipt of upstream data packets.

Data-triggered host creation does not rely on protocol triggers (DHCP, PPPoE) or management triggers (static hosts) to create each host, and is especially useful in the following cases:

- BNG manages subscribers under Layer 3 nodes (BRAS, CMTS, GGSN/PGW, and so on) and is not on the DHCP message path.
- BNG needs to manage large numbers of static hosts and bulk provisioning is required.

BNG authenticates, creates, and deletes subscriber hosts as follows:

Step 1. The subscriber SAP, including MSAP, receives a user packet that does not match existing anti-spoof table entries.

Step 2. BNG instantiates an IPoE session if there is no existing session with the same session key, and performs authentication using LUDB and RADIUS.

Note: Data-triggered ESM is supported only with IPoE sessions.

Step 3. A subscriber host is created with the ESM strings provided during authentication.

Step 4. The subscriber host is deleted when **session-timeout** or **idle-timeout** expires, CoA triggers a disconnect, SHCV check failure, or management (CLI, SNMP, etc) triggers a host deletion.

Note: There are no automatic triggers to delete a host if **session-timeout** or **idle-timeout** and SHCV are not configured.

9.17.1 Provisioning Data-triggered ESM

Data-triggered ESM can be enabled on a group interface. The following displays a sample configuration of data-triggered ESM:

```
subscriber-interface "SI1" create
```

```
group-interface "GI1" create
     arp-populate
     ipoe-session
          ipoe-session-policy "IS1"
     exit
     data-trigger
          no shutdown
     exit
```

An IPoE session and ARP population are mandatory when configuring data-triggered ESM.

The following packets can start data-trigger processing:

- IPv4 data packet
- IPv6 data packet
- ARP

To terminate IPv6 hosts that send neighbor RS/NS before sending data packets, **auto-reply** must be configured:

```
subscriber-interface "SI1" create
     group-interface "GI1" create
          ipv6
               auto-reply
                    neighbor-solicitation
               exit
```

For MSAP, the “data” trigger packet type can accept data triggers:

```
sap 1/1/2:200.* capture-sap create
     trigger-packet data
```

9.17.2 Authentication and Host Creation

Authentication of a data trigger can use LUDB configured in an IPoE session statement under a group interface.

To identify the source IPv4/IPv6 address of data-trigger packets, the IP prefix in the local user database can be configured with **host-identification**:

```
local-user-db "LUDB_DT" create
     ipoe
          match-list ip
          host "10.0.0.8" create
               host-identification
                    ip-prefix 10.0.0.8/29
               exit
          host "2001:1:b::1" create
```

```
        shutdown
        host-identification
             ip-prefix 2001:a:b::1/128
```

Note: Only one IP prefix can be configured for each host. A dual-stack host requires two local user database host entries if the IP prefix needs to be used for host identification.

For RADIUS authentication, the circuit ID includes the source IPv4/IPv6 address of the data-trigger packet:

```
authentication-policy "AUTH1" create
     user-name-format circuit-id
     include-radius-attribute
          circuit-id
     exit
```

If IPoE session policy uses circuit ID to identify each session, a new IPoE session is created for each source IPv4/IPv6 address. However, RADIUS can return the circuit ID to merge multiple IPoE sessions with the same SAP, MAC, and circuit ID into a single session.

A host is created using the IPv4/IPv6 source address of the data trigger (a /32 address for IPv4 or a /128 address for IPv6), but IPv6 data-triggered hosts can be created as an IPv6 prefix by configuring **ipv6-delegated-address** in the local user database host entry.

RADIUS can return the following AVPs to model the address/prefix of the data-triggered host:

- Framed-IP-Address: /32 IPv4 address of the host
- Framed-Route: managed IPv4 route with the host as next hop
- Alc-IPv6-Address: /128 IPv6 address of the host
- Delegated-IPv6-Prefix: IPv6 prefix of the host
- Framed-IPv6-Route: managed IPv6 route with the host as next hop

Information about multiple hosts can be returned in a single Access-Accept message when the **nh-mac anti-spoof** command is configured. This is mandatory when provisioning dual-stack hosts with the same SAP and MAC addresses with **nh-mac anti-spoof** configured, but is mutually exclusive with the CID key in the IPoE session policy.

9.17.3 DoS Protection

To authenticate data triggers, only the first packet is used for further processing. Subsequent packets from the same source are discarded until ESM host creation.

Data trigger packets are classified as all-unspecified protocol by Distributed CPU Protection (DCP).

9.17.4 DHCP Promotion

DHCP promotion allows data-triggered subscriber hosts to become DHCP hosts.

After a data-triggered host is created, DHCP packets sent by the client will start the DHCP promotion process as follows:

1. A DHCP Request/Renew/Rebind message comes from the data-triggered host.
2. Authentication using LUDB and RADIUS is performed. A RADIUS Access-Request message is sent if the authentication policy has re-authentication enabled. DHCP processing is done without authentication if re-authentication is not configured.
3. An Access-Accept message that contains ESM attributes is sent back from the AAA server. If an Access-Reject message is received, the data-triggered host is deleted.
4. A DHCP packet is relayed to the DHCP server.
5. A DHCP server replies with a DHCP Ack/Relay-Reply or Nak message.
6. When an Ack/Relay-Reply message is received, a lease state is created and the data-triggered host is promoted to the DHCP host.
7. If the DHCP server replies with Nak or with IP address information different from the existing data-triggered host, DHCP promotion will fail and the data-triggered host will be deleted.
8. With DHCP proxy, if the LUDB or AAA server returns IP information different from the data-triggered host, then DHCP promotion will fail and the data-triggered host will be deleted.
9. An interim accounting message is generated based on the configuration of the **radius-accounting-policy** command, as follows:
    1. **queue-instance-accounting**: interim-update with Alc-Acct-Triggered-Reason = Promotion of a Data-triggered host.
    2. **session-accounting**: interim-update with Alc-Acct-Triggered-Reason = Promotion of a Data-triggered host, if host-update is configured.

3. **host-accounting**: interim-update with Alc-Acct-Triggered-Reason = Promotion of a Data-triggered host

Data-triggered hosts can be promoted to DHCP proxy hosts by default. To promote data-triggered hosts using DHCP relay to an internal or external DHCP server, the Alc-Force-DHCP-Relay VSA will be included in Access-Accept messages to authenticate data-triggered hosts.

Note: DHCP relay promotion is only supported when using RADIUS. LUDB and NASreq is not supported.

Figure 151 **DHCP Promotion with DHCP Relay**

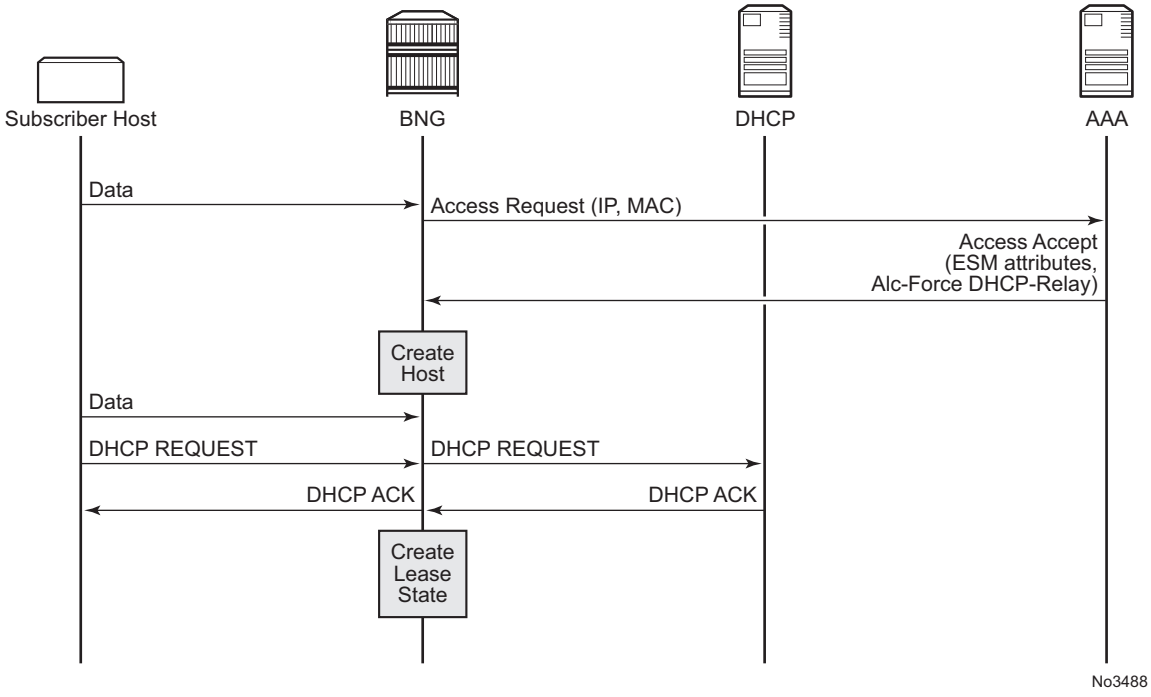

9.17.5 Data-Triggered SLAAC Hosts

A SLAAC host can be created instead of a data-triggered host on data-triggered authentication when the Framed-IPv6-Prefix is returned from the LUDB or AAA server and the IPv6 prefix in the AVP value matches the source address of the data-trigger.

Local address assignment (LAA) is not supported for data-triggered SLAAC hosts. Data-triggered subscriber management (**config**>**service**>**vprn**>**sub-if**>**grp-if**>**data-trigger**>**no shutdown**) and LAA commands (**config**>**service**>**vprn**>**sub-if**>**grp-if**> **lcl-addr-assign**>**ipv6**>**ipoe-wan | ipoe-slaac**) are mutually exclusive.

9.17.6 Stateful Multi-Chassis Redundancy (MCS)

Data-triggered subscriber hosts can be protected with stateful multi-chassis redundancy. Subscriber management MCS applications, under **config**>**redundancy**>**multi-chassis**>**peer**>**sync**>**sub-mgmt** also include data-triggered subscriber host information.

9.17.7 Stateless Multi-Chassis Redundancy

Stateless redundancy uses SRRP in the same form as the current implementation for master/standby selection and peer liveness detection, but does not need subscriber state synchronization using MCS that requires CPM/IOM resources to be on standby node.

Stateless redundancy has the following characteristics:

- IPoE sessions and their hosts (DHCP,DHCP6,SLAAC, and DT) are not synchronized via the subscriber management IPoE MCS client.
- SRRP is synchronized over MCS.
- BNG in the SRRP master state only processes the data-trigger, and authenticates and creates a subscriber host state.
- BNG in the SRRP state other than master, discards all upstream packets.
- Shunt and redundant interfaces are not supported.
- After an SRRP switchover, a new master BNG starts processing subscriber traffic. The old master BNG deletes all IPoE sessions and IPoE subscriber hosts on an SRRP switchover. Accounting stops will be sent to indicate SRRP mastership changed to the backup BNG node. Accounting messages are different based on the **radius-accounting-policy** configuration.
  1. **queue-instance-accounting**: accounting-stop with Alc-Error-Code = Node has switched to stateless backup
  2. **session-accounting**: interim-update with Alc-Acct-Triggered-Reason = Node has switched to stateless backup, upon each stack deletion if host-update is configured, and accounting-stop with Alc-Error-Code = Node has switched to stateless backup

3. **accounting**: accounting-stop with Alc-Error-Code = Node has switched to stateless backup, upon each host deletion

- The DHCP local server state can be synchronized for DHCP promotion.
- DHCP promotion for BRG requires lease state synchronization between redundant BNGs, which disables protocol-triggered IPoE ESM without an IPoE session. This requires the **config**>**service**>**ies** | **vprn**>**sub-if**>**grp-if**>**ipoe-session**>**stateless-redundancy** command to be configured.
  - host-PPPoE (PTA and LAC) processing can be tied with an SRRP instance for stateless redundancy. Only the master BNG processes PPPoE/PPP control plane data. A non-master BNG does not send LCP echo messages. After switchover, subscriber hosts will retry to connect after echo timeout.
- Supports both static SAP and MSAP.

Figure 152 **Figure: Stateless Multi-Chassis Redundancy**

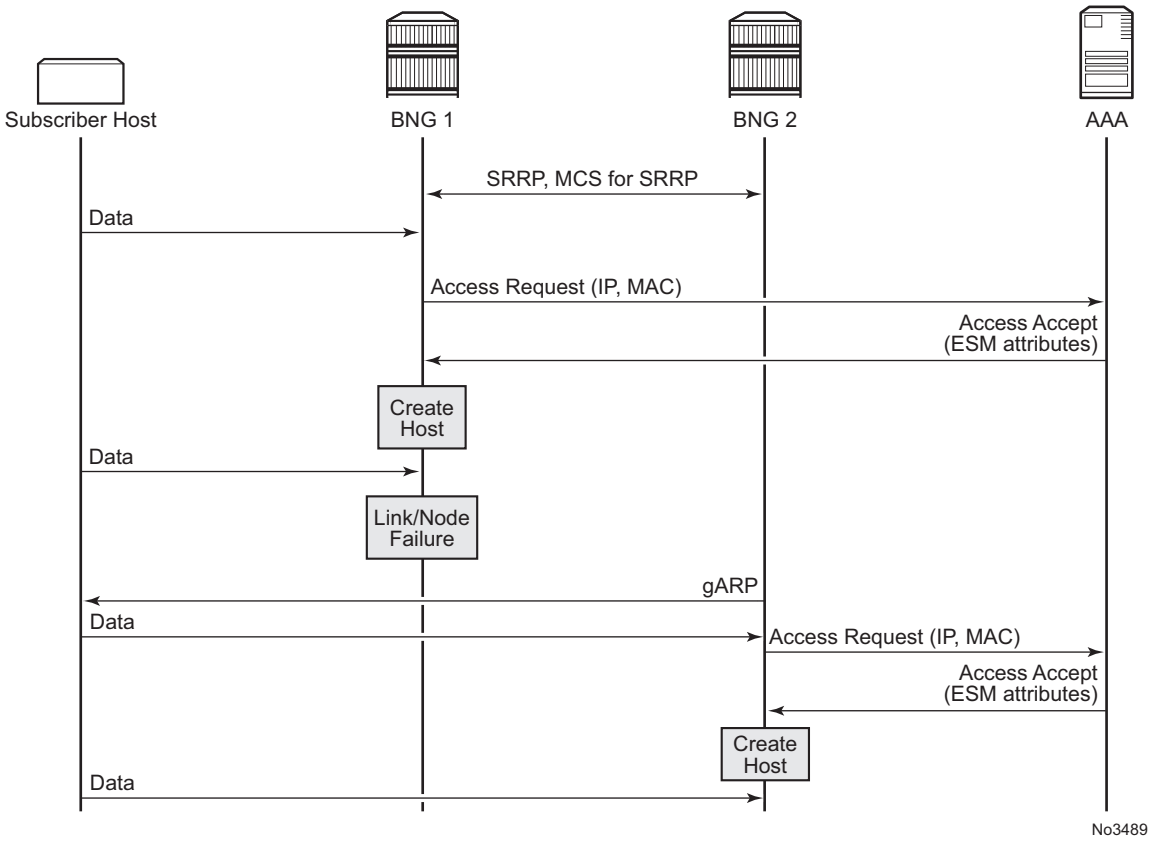

9.17.7.1 MSAP Support

Stateless redundancy does not have information on MSAP because MCS is not used for synchronizing subscriber host information.

Static SAPs need to be configured to rewrite FDBs on Layer 2 switches via G-ARP in access/aggregation networks.

to rewrite FDBs via G-ARP

```
group-interface <ip-int-name>
sap <sap-id> create
no shutdown
```

This requires the aggregation network to re-learn MAC for multiple C-VLANs using a single G-ARP from a specific C-VLAN.

Two typical scenarios are supported:

- MSAP has qinq encap, and aggregation switches have per-SVLAN FDB
- MSAP has dot1q encap, and aggregation switches have per-SVLAN FDB (Aggregation switches push SVID)

9.18 RADIUS Subscriber Services

RADIUS subscriber services enable the activation and deactivation of subscriber functions via RADIUS Access-Accept or CoA messages. Each subscriber service can have its own RADIUS accounting session.

The subscriber service functionality is built using the flexible RADIUS Python script interface to populate the subscriber service data structure using a parameter list received in subscriber service-specific RADIUS Vendor Specific Attributes (VSAs). The format and content of the VSA parameter list is defined by the operator. An accounting start/stop is sent when the subscriber service is activated/deactivated. Optionally, interim updates can be sent in intervals that can be specified per subscriber service instance. Accounting interim updates and stop messages contain the subscriber service-related statistics (time or volume and time).

Subscriber services can be activated on a single-stack or dual-stack PPPoE or IPoE session or on a single-stack IPv4 host.

Subscriber service functionality can be built with:

- QoS overrides: changing queue or policer parameters (PIR/CIR rates and CBS/MBS burst sizes), adapting rates of a parent scheduler, root arbiter, or subscriber aggregate rate
- PCC rules: applying QoS or filter actions to a set of IP flows

9.18.1 Subscriber Service Building Blocks

Figure 153 shows the building blocks required to activate or deactivate a subscriber service.

Figure 153 **Subscriber Services Building Blocks**

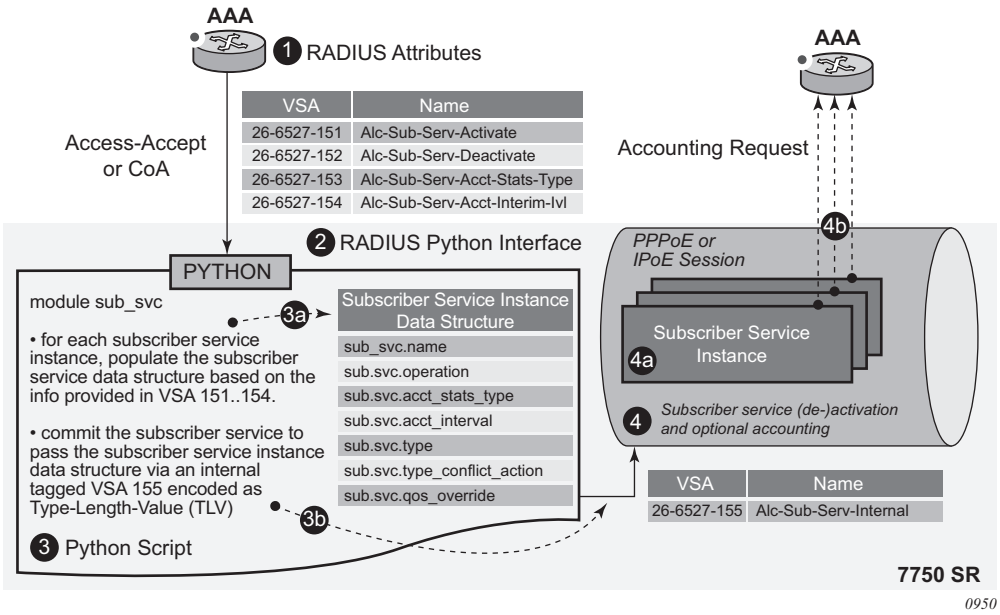

Each of the building blocks is described in the following sections.

9.18.1.1 RADIUS Access-Accept or CoA Message with Subscriber Service Activate or Deactivate VSAs

A subscriber service instance is activated from the RADIUS server via an Access-Accept or CoA message for a PPPoE or IPoE session. Deactivation of a subscriber service instance can be achieved via a RADIUS CoA message or is implicit when the associated subscriber session terminates.

To activate a subscriber service, the Alc-Sub-Serv-Activate (VSA) is used, and to deactivate a subscriber service, the Alc-Sub-Serv-Deactivate VSA is used. The formats of the Alc-Sub-Serv-Activate and Alc-Sub-Serv-Deactivate VSAs can be freely defined by the operator as long as they match with the Python script that is used to commit the subscriber service instance activation or deactivation.

For example, to change the upstream and downstream bandwidth of an IPoE session, the following format can be defined:

```
rate-limit;<upstream_bw_in_kbps>;<downstream_bw_in_kbps>
```

To activate a subscriber service with an upstream bandwidth of 5 Mb/s and a downstream rate of 50 Mb/s, the following VSA can then be included in a RADIUS Access-Accept or CoA message:

```
Alc-Sub-Serv-Activate = "rate-limit;5120;30720"
```

To deactivate the same subscriber service and revert back to the initial bandwidth, the following VSA can be included in a RADIUS CoA message:

```
Alc-Sub-Serv-Deactivate = "rate-limit;5120;30720"
```

To deactivate a subscriber service instance, its unique name must be used. In the example above, the name equals “rate-limit;5120;30720”.

To start an accounting session when the subscriber service instance is activated, the following attributes can be included in the Access-Accept or CoA message:

```
Alc-Sub-Serv-Acct-Stats-Type = volume-time | time
Alc-Sub-Serv-Acct-Interim-Ivl = <update-interval>
```

For example, the Alc-Sub-Serv-Acct-Stats-Type attribute value is set to “volume-time” to include both the session time for time-based billing and standard counters for volume statistics collection. The Alc-Sub-Serv-Acct-Interim-Ivl attribute sets the interval for interim updates of the subscriber service instance accounting.

See the Subscriber Services RADIUS VSAs section for details on RADIUS attributes.

See the Subscriber Service RADIUS Accounting section for details on subscriber service instance accounting.

9.18.1.2 RADIUS Python Interface

The SR OS RADIUS Python interface is used to interpret the parameters specified in the subscriber service-specific VSAs and to generate an internal proprietary format VSA representing the subscriber service instance to activate or deactivate, as shown in Figure 154.

Figure 154 **RADIUS Python Interface**

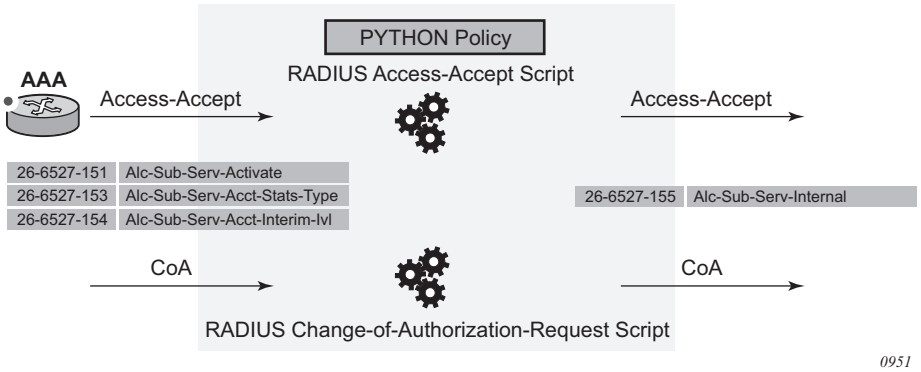

The SR OS only requires the Alc-Sub-Serv-Internal VSA [26-6527-155] to activate or deactivate the subscriber service. Subscriber service-specific VSAs [26-6527-151..154] are available to be used in the subscriber service Python script but are ignored in the SR OS.

See the Python Script section for details on the subscriber services Python script functions and operation.

A Python script must be configured for RADIUS Access-Accept and CoA messages; for example:

```
config
    python
        python-script "subsvc-1" create
            primary-url "ftp://user:pwd@10.1.1.1/./py/subsvc-1.py"
            no shutdown
        exit
        python-policy "py-policy-subsvc-1" create
            radius access-accept direction ingress script "subsvc-1"
            radius change-of-authorization-request direction ingress script "subsvc-
1"
        exit
    exit
```

The Python policy must then be applied to the **radius-server-policy** to pass the Access-Accept messages to the Python script and to the RADIUS server to pass the CoA messages to the Python script; for example:

```
configure
    router
        radius-server
            server "server-1" address 10.1.1.2 secret <secret> create
                accept-coa
                python-policy "py-policy-subsvc-1"
            exit
        exit
```

```
exit
aaa
    radius-server-policy "aaa-server-policy-1" create
        python-policy "py-policy-subsvc-1"
        servers
            access-algorithm round-robin
            router "Base"
            server 1 name "server-1"
        exit
    exit
exit
```

9.18.1.3 Python Script

The RADIUS Access-Accept and CoA messages are passed to the configured RADIUS Python scripts. As shown in Figure 155, the function of the subscriber service Python script is to interpret the subscriber service-specific VSAs that contain the subscriber service instance parameters and to generate a new Alc-Sub-Serv-Internal VSA containing the information required to activate the actual subscriber service on the PPPoE or IPoE session.

This section covers the basics to understand the functionality of a subscriber service Python script. See the Subscriber Services Python API section for a detailed description of the alc.sub_svc Python module containing functions and data structures used to define and activate a subscriber service instance.

Figure 155 **Subscriber Service Python Script Operation**

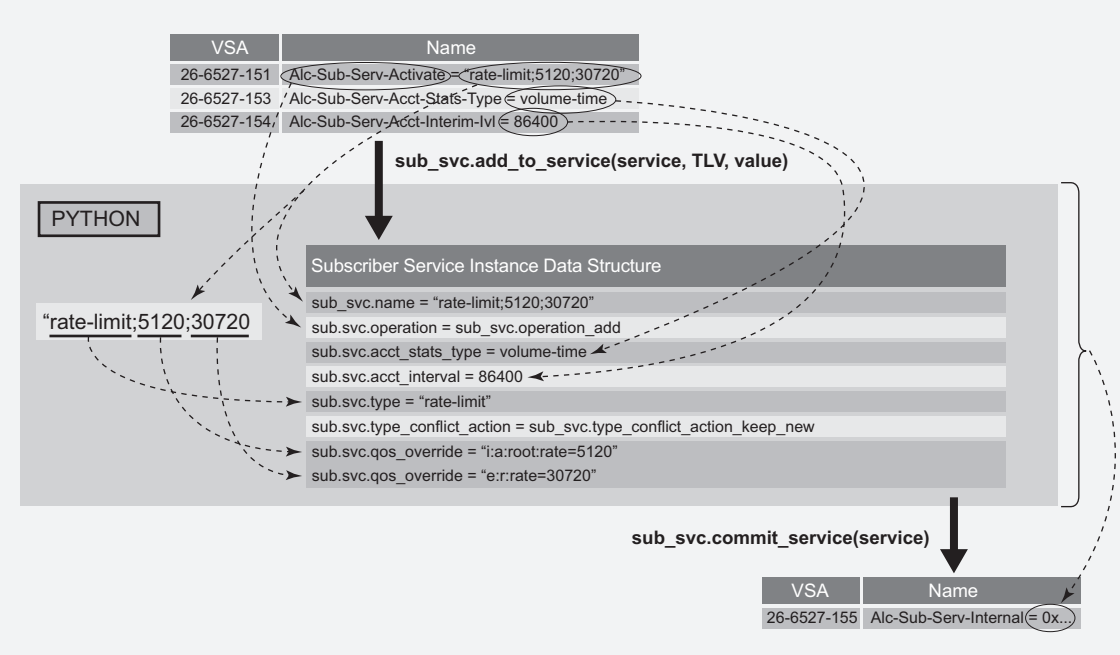

The alc.sub_svc Python module contains the required functions and data structure to commit a subscriber service, including:

- TLVs to build a data structure describing the subscriber service functionality, such as QoS overrides or PCC rules
- functions to populate the data structure such as sub_svc.add_to_service(), sub_svc.pccrule.add_to_pccrule(), and sub_svc.flow.add_to_flow()
- a function to commit the subscriber service and create the internal VSA from the data structure: sub_svc.commit_service()

9.18.1.3.1 Python Script Example

In this section an example of a Python script is explained that enables the activation or deactivation of a subscriber service.

In the example, it is assumed that only a single subscriber service is activated and/or deactivated per RADIUS message (no tagged VSAs are used) and that only a single Alc-Sub-Serv-Activate or Alc-Sub-Serv-Deactivate VSA is present (no concatenation of VSAs is required).

To change the upstream root arbiter rate and downstream aggregate rate bandwidth of an IPoE session, send the following parameters in the subscriber service activate VSA:

```
rate-limit;<upstream_bw_in_kbps>;<downstream_bw_in_kbps>
```

During the bandwidth change, the traffic should be accounted for in a separate accounting session.

To import the required modules in the Python script:

```
# Python - imports
import struct
from alc import radius
from alc import sub_svc
```

The alc.radius module provides the API access to the RADIUS VSAs in Access-Accept and CoA messages.

The alc.sub_svc module allows the API to activate and deactivate subscriber services.

The struct module is a Python module used in the example to convert data obtained from the RADIUS API as a string into Python integer values.

The following constants are used in the script:

```
# Python - constants
# VSA vendor ID
ALC = 6527
# ALC Radius VSA
SUB_SERVICE_ACTIVATE = 151
SUB_SERVICE_DEACTIVATE = 152
SUB_SERVICE_ACCT_STATS_TYPE = 153
SUB_SERVICE_ACCT_INTERIM_IVL = 154
```

The main flow in a subscriber service Python script is to first process the subscriber service deactivations, followed by the subscriber service activations. Optionally, the subscriber service-specific VSAs can be removed from the RADIUS message as they are not required for further processing in the SR OS:

```
# Python - main()
deactivate_services()
activate_services()
radius.attributes.clearVSA(ALC, SUB_SERVICE_ACTIVATE)
radius.attributes.clearVSA(ALC, SUB_SERVICE_DEACTIVATE)
radius.attributes.clearVSA(ALC, SUB_SERVICE_ACCT_STATS_TYPE)
radius.attributes.clearVSA(ALC, SUB_SERVICE_ACCT_INTERIM_IVL)
```

The function to deactivate a subscriber service executes the following steps:

- checks if an Alc-Sub-Serv-Deactivate VSA [26-6527-152] is present
- if present:
  - builds the subscriber service data structure. For a subscriber service deactivation, only the subscriber service instance name is required.
  - commits the subscriber service deactivation

```
# Python - deactivate_services()
def deactivate_services():
  value = radius.attributes.getVSA(ALC, SUB_SERVICE_DEACTIVATE)
  if value != '':
    service = []
    sub_svc.add_to_service(service, sub_svc.name, value);
    sub_svc.add_to_service(service, sub_svc.operation, sub_svc.operation_del);
    sub_svc.commit_service(service)
```

The function to activate a subscriber service executes the following steps:

- checks if an Alc-Sub-Serv-Activate VSA [26-6527-151] is present
- if present:
  - separates the parameters in the attribute value with a semicolon as the delimiter. The parameters are the subscriber service type, the upstream bandwidth in kb/s, and the downstream bandwidth in kb/s.
  - If the subscriber service type is rate-limit, builds the subscriber service data structure:
    - specifies that the service will be added (activate)
    - adds the complete attribute value as the subscriber service instance name. This name should also be used to deactivate the subscriber service.
    - sets the subscriber service type and conflict action. With the conflict action set to “keep new”, if a new subscriber service of the same type is activated, the old one is deactivated first.
    - configures the upstream and downstream bandwidth with a QoS override on the ingress root arbiter and egress aggregate rate
  - If the subscriber service type is something else (an unknown service type), a warning is printed for debugging purposes. No action is performed in this case.
  - If an Alc-Sub-Serv-Acct-Stats-Type VSA [26-6527-153] is present, adds the corresponding stats-type to the subscriber service data structure. The function also checks whether a Alc-Sub-Acct-Serv-Interim-Ivl VSA [26-6527-154] is present and adds the interval to the subscriber service data structure.

Both the stats-type and interim interval must be specified as integers in the subscriber service data structure. As the RADIUS API returns an octet string, a conversion is required. The struct.unpack() function is used for this purpose.

– commits the subscriber service activation

```
# Python - activate_services()
def activate_services():
  # Subscriber Service Activate VSA
  value = radius.attributes.getVSA(ALC, SUB_SERVICE_ACTIVATE)
  if value != '':
    values = value.split(';')
    if values[0] == "rate-limit":
      service = []
      sub_svc.add_to_service(service, sub_svc.operation, sub_svc.operation_add);
      sub_svc.add_to_service(service, sub_svc.name, value);
      sub_svc.add_to_service(service, sub_svc.type, 'rate-limit')
      sub_svc.add_to_service(service, sub_svc.type_conflict_action, sub_svc.type_con
flict_action_keep_new)
      sub_svc.add_to_service(service, sub_svc.qos_override, 'i:a:root:rate=' +
values[1])
      sub_svc.add_to_service(service, sub_svc.qos_override, 'e:r:rate=' + values[2])
    else:
      print "WARNING - Unknown service type :", values[0]
      return
    # Subscriber Service Accounting VSA
    stats_type = radius.attributes.getVSA(ALC, SUB_SERVICE_ACCT_STATS_TYPE)
    if stats_type != '':
      sub_svc.add_to_service(service, sub_svc.acct_stats_type, struct.unpack("!I", s
tats_type)[0]);
      interval = radius.attributes.getVSA(ALC, SUB_SERVICE_ACCT_INTERIM_IVL)
      if interval != '':
        sub_svc.add_to_service(service, sub_svc.acct_interval, struct.unpack("!I", i
nterval)[0]);
    # Activate the Subscriber Service
    sub_svc.commit_service(service)
```

The result of the sub_svc.commit_service() function is an Alc-Sub-Serv-Internal VSA [26-6527-155] that contains the required data for the SR OS to activate the corresponding subscriber services.

9.18.1.4 Subscriber Service Instance Activation or Deactivation with Optional RADIUS Accounting

When the SR OS receives an Alc-Sub-Serv-Internal VSA [26-6527-155] in an Access-Accept or CoA message as a result of the Python script sub_svc.commit_service() function, the corresponding subscriber services are activated and/or deactivated. Optionally, an accounting session can be started for each subscriber service instance.

The following is an example of the **show service sub-services** command.

```
# show service sub-services
===============================================================================
Subscriber service table
===============================================================================
Username        : cpe2@domain1.com
Subscriber      : cpe2@domain1.com                      Type     : PPP
SAP             : 1/1/4:1201.5                          Service  : 1000
MAC Address     : 00:51:00:00:01:01                     PPPoE-SID: 114
IP              : 10.1.1.3
Interface-ID    : 02:51:00:FF:FE:00:01:01
-------------------------------------------------------------------------------
  Service instance : rate-limit;5120;30720
  up time          : 0d 00:01:06
  type             : rate-limit
  acct sess id     : 144DFF00000A2556B757DA
  multi sess id    : 144DFF00000A1156B71E44
  acct type        : volume-time
  acct ivl         : 1d 00:00:00
  input octets     : 0
  input packets    : 0
  output octets    : 0
  output packets   : 0
  actions          :
  QoS-override
    ingress arbiter "root" rate 5120
  QoS-override
    egress aggregate rate limit 30720
-------------------------------------------------------------------------------
number of subscriber services found: 1
===============================================================================
```

The following rules apply for a subscriber service instance.

- A subscriber service instance can be activated on a:
  - single-stack or dual-stack PPPoE session
  - single-stack or dual-stack IPoE session
  - single-stack IPoEv4 host
- A subscriber service instance is uniquely identified by its name. To deactivate a subscriber service instance, the subscriber service name must be referenced.
- A subscriber service cannot be modified. To update a subscriber service instance, the old subscriber service must be deactivated and a new one activated.
- Multiple subscriber services can be activated or deactivated in a single Access-Accept or CoA message. Tagged subscriber service-specific VSAs are available for this purpose.
- A single subscriber service instance can have multiple actions, such as multiple QoS overrides and/or multiple PCC rules.

- The conflict action determines the behavior when multiple subscriber services of the same type are activated for the same PPPoE or IPoE session:
  - Conflict action = none

    Multiple instances of the same type can be activated.
  - Conflict action = keep new

    Only a single subscriber service instance of the same type is allowed.

    When a new subscriber service instance of the same type is activated on a single PPPoE or IPoE session, the old instance is deactivated and the new subscriber service instance is activated.
  - Conflict action = keep old

    Only a single subscriber service instance of the same type is allowed.

    When a subscriber service instance of the same type is activated on a single PPPoE or IPoE session, the new subscriber service instance is rejected.
- Multi-Chassis Synchronization (MCS) is not supported for subscriber services.

9.18.2 Subscriber Services RADIUS VSAs

Subscriber services VSAs 26-6527-151..154 are used as inputs for the subscriber service Python script and are ignored by the SR OS.

The subscriber service VSAs can be tagged to allow activation and deactivation of multiple subscriber service instances with a single RADIUS Access-Accept or CoA message.

Table 90 lists the subscriber services RADIUS VSAs. See the RADIUS Attributes Reference Guide for a complete description of the subscriber services VSAs.

Table 90 **Subscriber Services RADIUS VSAs**

Attribute ID	Attribute Name	Description
26-6527-151	Alc-Sub-Serv-Activate (string)	Activate a subscriber service. The attribute contains parameters as input for the subscriber service Python script to define and activate a subscriber service. Multiple VSAs can be included per message if the parameter list does not fit in a single attribute.

Table 90 **Subscriber Services RADIUS VSAs (Continued)**

Attribute ID	Attribute Name	Description
26-6527-152	Alc-Sub-Serv-Deactivate (string)	Deactivate a subscriber service. The attribute contains parameters as input for the subscriber service Python script to deactivate a subscriber service. Multiple VSAs can be included per message if the parameter list does not fit in a single attribute.
26-6527-153	Alc-Sub-Serv-Acct-Stats-Type (integer)	Enable or disable subscriber service accounting and specify the statistics type: volume and time or time only. Values: 1=off, 2=volume-time and 3=time
26-6527-154	Alc-Sub-Serv-Acct-Interim-Ivl (integer)	The interim accounting interval in seconds at which Acct-Interim-Update messages should be generated for subscriber service accounting. A value of 0 (zero) corresponds to no interim update messages. The maximum value is 300 seconds (values 1 to 300).
26-6527-155	Alc-Sub-Serv-Internal	For internal use only. Its value is the result of the subscriber service commit function in Python. (sub_svc.commit_service).

9.18.3 Subscriber Service RADIUS Accounting

The configuration for subscriber service instance accounting sessions is obtained from the RADIUS accounting policies configured in the subscriber profile of the parent subscriber:

- The accounting destinations: RADIUS server configuration

  Subscriber service instance accounting is sent to multiple destinations when multiple RADIUS accounting policies are configured (for example, **config>subscr-mgmt>sub-prof# radius-accounting policy acct-policy-1 acct-policy-2**)

  Up to five accounting policies can be configured. Each policy is independent of the other, with its own accounting mode, update interval, and include attributes. Since resources are limited for sending out RADIUS account messages, contact your Nokia technical support representative for recommendations.
- The accounting attributes (**include-radius-attribute**) with some exceptions.

Standard accounting attributes are always used for volume accounting, regardless of the configuration in the RADIUS accounting policy.

An untagged [26-6527-151] Alc-Sub-Serv-Activate attribute is included in all subscriber service instance accounting messages. Its value is the subscriber service instance name (alc.sub_svc.name).

- Accounting interim updates for subscriber service instance accounting are controlled independently from the parent subscriber accounting session:
  - Enabling accounting interim updates for a subscriber service instance is achieved by setting the sub_svc.acct_interval TLV in the subscriber services Python script to a value greater than 0: sub_svc.add_to_service(service, sub_svc.acct_interval, 3600)
  - Disabling accounting interim updates for a subscriber service instance is achieved by setting the sub_svc.acct_interval TLV in the subscriber services Python script to 0: sub_svc.add_to_service(service, sub_svc.acct_interval, 0)

If the parent subscriber has no RADIUS accounting policy configured, subscriber service instance accounting cannot be enabled.

If the parent subscriber has a RADIUS accounting policy configured, subscriber service instance accounting can be disabled by setting the sub_svc.acct_stats_type TLV in the subscriber service Python script to 1 (Off):

```
sub_svc.add_to_service(service, sub_svc.acct_stats_type, 1)
```

Each subscriber service instance has a unique accounting session ID included as [44] Acct-Session-Id. The [50] Acct-Multi-Session-Id contains the accounting session ID of the parent PPPoE or IPoE session as shown in Table 91:

Table 91 **Subscriber Service Accounting Multi-Session ID**

Subscriber Service Parent	[50] Acct-Multi-Session-Id
PPPoE session	Session Acct-Session-Id of the PPPoE session: `show service id <service-id> pppoe session detail`
IPoE session	Session Acct-Session-Id of the IPoE session: `show service id <service-id> ipoe session detail`
IPoEv4 host	Host Acct-Session-Id of the IPoEv4 host: `show service id <service-id> subscriber-hosts detail`

The content of the volume counters when the subscriber service accounting statistics type equals **volume-time** is determined by the subscriber service action. For details, see the subscriber service sections that follow.

9.18.4 Accounting-Only Subscriber Service

An accounting-only subscriber service has no specific action such as **qos-override** or **pccrule** defined and has subscriber service instance accounting enabled.

An example of when an accounting-only subscriber service would be used is if additional accounting data is needed for a specific time period in the lifetime of a PPPoE or IPoE session.

The sub_svc.acct_stats_type TLV in the subscriber services Python script must be set to a value different from 1 (Off) to enable subscriber service instance accounting. For example:

```
# Python example - Accounting only subscriber service
service = []
sub_svc.add_to_service(service, sub_svc.operation, sub_svc.operation_add)
sub_svc.add_to_service(service, sub_svc.name, 'subsvc-acct-1')
sub_svc.add_to_service(service, sub_svc.type, 'acct-only')
sub_svc.add_to_service(service, sub_svc.type_conflict_action, sub_svc.type_conflict_
action_none)
sub_svc.add_to_service(service, sub_svc.acct_stats_type, 2)
sub_svc.add_to_service(service, sub_svc.acct_interval, 3600)
sub_svc.commit_service(service)
```

The volume counters for subscriber service statistics type **volume-time** contain the aggregate forwarded octets and packets of the parent PPPoE or IPoE session **sla-profile** instance since the start of the subscriber service.

9.18.5 QoS Override-based Subscriber Service

A subscriber service instance with a **qos-override** action overrides queue/policer parameters (CIR, PIR, CBS, MBS) configured at the **sla-profile** level and hierarchical QoS parameters (aggregate rate, scheduler rate, or root/intermediate arbiter rate) configured at the **sub-profile** level.

An example of when a QoS override-based subscriber service would be used is to temporarily offer higher bandwidth and charge for the volume consumed during this period. Figure 156 shows the volume statistics that are reported for the PPPoE or IPoE session accounting and for each of the subscriber service instances accounting sessions that were activated and deactivated.

Figure 156 **Sample Use Case: QoS Override-based Subscriber Services**

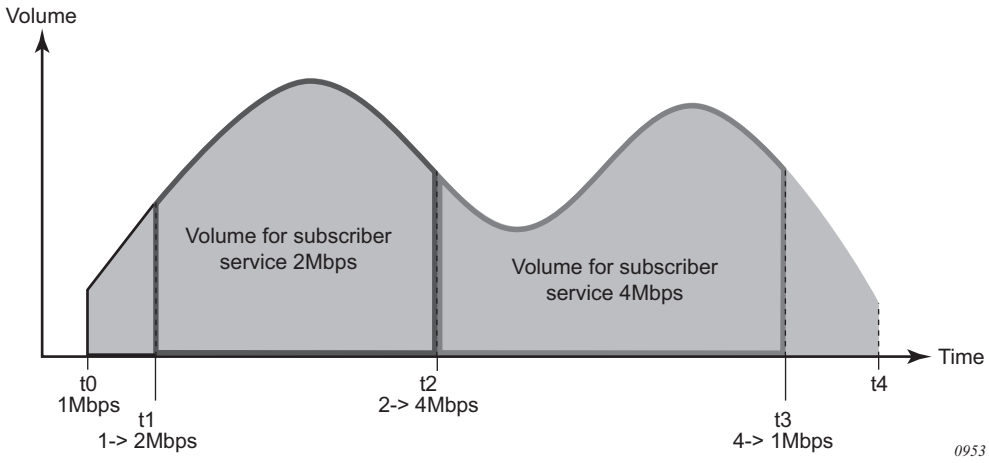

- At t0, a subscriber connects and starts consuming bandwidth at the subscriber base egress rate of 1 Mb/s. An accounting session is started for the subscriber.
- At t1, a QoS override-based subscriber service instance is activated that boosts the download bandwidth to 2 Mb/s. A 2 Mb/s QoS override-based subscriber service instance accounting session is started.
- At t2, the 2 Mb/s QoS override-based subscriber service instance is deactivated and replaced with a new one that increases the download bandwidth to 4 Mb/s. The 2 Mb/s QoS override-based subscriber service accounting session is terminated, reporting the volume consumed during t1 to t2. A new 4 Mb/s QoS override-based subscriber service accounting session is started.
- At t3, the 4 Mb/s QoS override-based subscriber service instance is deactivated. The download bandwidth is reduced to the 1 Mb/s base rate again. The 4 Mb/s QoS override-based subscriber service accounting session is terminated, reporting the volume consumed during t2 to t3.
- At t4, the subscriber disconnects. The accounting session for the subscriber is terminated, reporting the volume consumed during t0 to t4.

The sub_svc.qos_override TLV in the subscriber services Python script adds a **qos-override** action. For example:

```
# Python example - QoS override subscriber service
service = []
sub_svc.add_to_service(service, sub_svc.operation, sub_svc.operation_add)
sub_svc.add_to_service(service, sub_svc.name, 'subsvc-qos-override-1')
sub_svc.add_to_service(service, sub_svc.type, 'qos-override')
sub_svc.add_to_service(service, sub_svc.type_conflict_action, sub_svc.type_conflict_
action_keep_new)
sub_svc.add_to_service(service, sub_svc.qos_override, 'e:q:1:pir=20000,cir=0')
sub_svc.add_to_service(service, sub_svc.qos_override, 'e:q:2:pir=5000,cir=5000')
sub_svc.add_to_service(service, sub_svc.qos_override, 'i:p:1:pir=1000,cir=0')
```

```
sub_svc.add_to_service(service, sub_svc.qos_override, 'e:r:rate=20000')
sub_svc.commit_service(service)
```

Actual values are used to populate the subscriber service data structure in this example; typically, these values are sent as parameters in subscriber service-specific VSAs.

Although a QoS override-based subscriber service instance is activated for a PPPoE or IPoE session, the overrides are applied at the SLA profile instance and/or subscriber level.

QoS override-based subscriber services have precedence over dynamic QoS overrides (RADIUS Alc-Subscriber-QoS-Override VSA or Gx Charging-Rule-Definition with QoS-Information) on the PPPoE or IPoE session.

- When a dynamic QoS override such as via the Alc-Subscriber-QoS-Override VSA is active on the parent PPPoE or IPoE session and a QoS override-based subscriber service is activated, the action from the QoS override-based subscriber service is installed. When the QoS override-based subscriber service is later deactivated, the original dynamic QoS override is restored.
- When a dynamic QoS override such as via the Alc-Subscriber-QoS-Override VSA is received for a parent PPPoE or IPoE session with an active QoS override-based subscriber service, the QoS override-based subscriber service remains installed. If the QoS override-based subscriber service is later deactivated, the previously received dynamic QoS override is installed.
- When both a QoS override-based subscriber service activation and a dynamic QoS override such as via the Alc-Subscriber-QoS-Override VSA is received in a single message for the parent PPPoE or IPoE session, the QoS override-based subscriber service action is installed and the dynamic QoS override is stored for later reference. When the QoS subscriber service is later deactivated, the previously received dynamic QoS override is installed.

The installed QoS override actions can be verified in the output of the **show service active-subscribers detail** CLI command.

The volume counters for subscriber service statistics type **volume-time** contain the aggregate forwarded octets and packets of the parent PPPoE or IPoE session **sla-profile** instance since the start of the subscriber service.

QoS override-based subscriber services are stored in the subscriber-mgmt persistency file.

9.18.6 PCC Rule-based Subscriber Services

Policy and Charging Control (PCC) rules are defined in the 3GPP PCC architecture and used in the Diameter Gx application as a collection of parameters that enable IP traffic flows to be identified, QoS parameters and filtering actions to be applied to these flows, and charging to be performed on them. The use of PCC rules in policy management via the Diameter Gx interface is described in the PCC Rules.

The same PCC rule construct is used in RADIUS subscriber services to enable IP flow-based actions and accounting.

- Identify IP flows based on 5-tuple (protocol, source and destination IP address, source and destination port) and/or DSCP value.
- Apply QoS actions to these flows, such as rate limiting or change of forwarding class.
- Apply filter actions to these flows, such as forward, drop, HTTP redirect, redirect to a next hop or routing instance.
- Enable RADIUS accounting to report the forwarded octets and packets of the IP flows.

IP flow-based accounting is one use for a subscriber service using PCC rules. Activated via a self-service portal or as part of an Internet subscription package, applications identified by a 5-tuple receive specific treatment, such as bandwidth increase, expedited forwarding, or zero rating. Volume and time statistics for the application data are available in the subscriber service RADIUS accounting session. This is shown in Figure 157.

Figure 157 **Example: PCC Rule-based Subscriber Service**

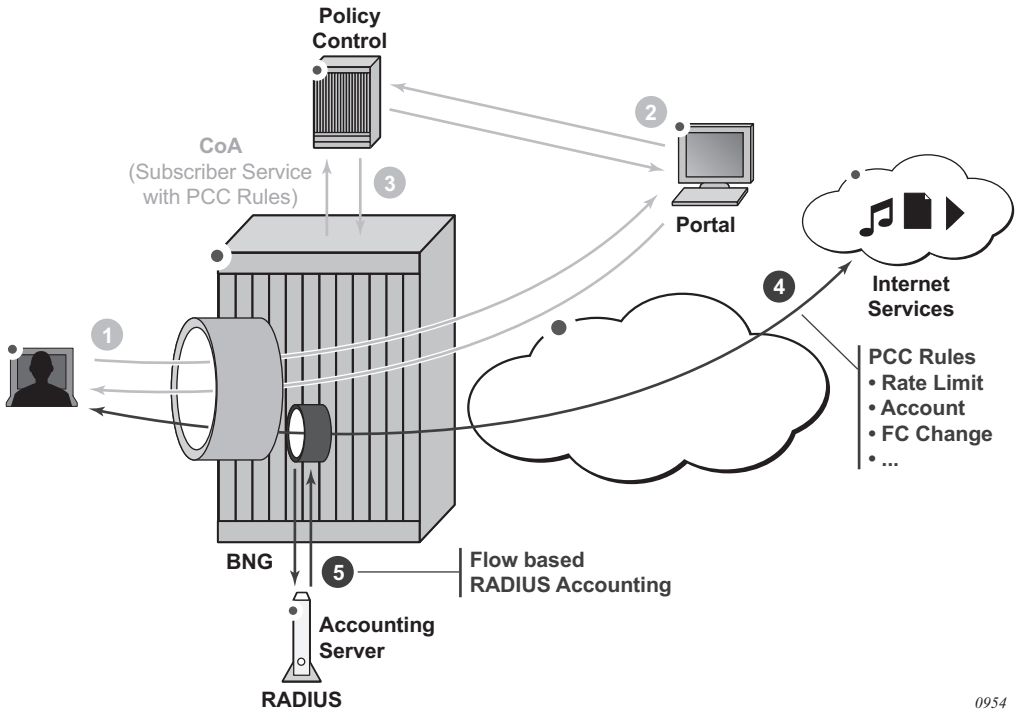

1. The user subscribes to a service via a web portal.
2. The policy control is informed of the new subscription.
3. The policy control instructs the BNG via RADIUS to activate a subscriber service with PCC rules.
4. The subscriber service is instantiated on the BNG. A dynamic policer is spawned to optionally rate-limit and/or count the application traffic.
5. The subscriber service instance accounting session reports the volume of traffic forwarded as part of this service.

9.18.6.1 PCC Rule Actions

A PCC rule is a unidirectional set of IP flows sharing a same set of actions. IPv4 and IPv6 flows can be combined within the same PCC rule.

A PCC rule name must be unique for each rule applied on a single PPPoE or IPoE session. For optimal PCC rule sharing, it is recommended that the same PCC rule name be used when its content is the same (that is, the same set of flows and same set of actions).

An IP flow is identified by a combination of:

- 5-tuple (IPv4 or IPv6): src-ip, src-port, dst-ip, dst-ip, protocol/next-header
- DSCP value

The CLI equivalent is:

```
match protocol | next-header <protocol>
       src-ip <ip-address>
       dst-ip <ip-address>
       src-port eq <port> | range <start-port> <end-port>
       dst-port eq <port> | range <start-port> <end-port>
       dscp <dscp>
   exit
```

Supported actions include: forward, drop, redirect to ip next hop, redirect to a routing instance, HTTP redirect, forwarding class change, rate-limit, and account.

With a specified set of actions, PCC rules are instantiated in the SR OS via IP criteria or IPv6 criteria entries in SAP ingress or SAP egress QoS polices and/or in IP or IPv6 filter entries. A PCC rule precedence value determines the relative order of different PCC rules when inserted in the QoS or filter policy: a rule with a lower precedence value will be applied before a rule with a higher precedence value. Rules with the same precedence can be automatically optimized; the relative order in which they are applied is determined by the system for optimal sharing. Rules with no precedence are applied at the end and are also automatically optimized.

Table 92 and Table 93 provide an overview of the PCC rule actions and where they apply.

Table 92 **Subscriber Service PCC Rule Actions Resulting in QoS Policy Changes**

Action	Direction	Description

Forwarding Class (FC) change	Ingress/Egress	changes the QoS forwarding class
		CLI equivalent:
		`config>qos`
		`  sap-ingress | sap-egress <id> create`
		`    ip-criteria | ipv6-criteria`
		`      entry <id> create`
		`        match`
		`          <5-tuple | dscp>`
		`        exit`
		`        action fc <fc>`
		`      exit`
		`    exit`

Table 92 **Subscriber Service PCC Rule Actions Resulting in QoS Policy Changes (Continued)**

Action	Direction	Description

Rate-limit (PIR/CIR)	Ingress/Egress	rate-limit traffic matching the specified flows
		• spawn a dynamic policer per direction (if not already present)
		• set the PIR/CIR value of the dynamic policer
		• map the flows to the policer
		The forwarded octets and packets statistics of the dynamic policer are included in subscriber service accounting.
		CLI equivalent:
		`config>qos`
		`  sap-ingress | sap-egress <id> create`
		`    policer 1  # dynamic policer`
		`      rate <pir> cir <cir>`
		`    exit`
		`    ip-criteria | ipv6-criteria`
		`      entry <id> create`
		`        match`
		`          <5-tuple | dscp>`
		`        exit`
		`        action policer 1`
		`      exit`
		`    exit`

Table 92 **Subscriber Service PCC Rule Actions Resulting in QoS Policy Changes (Continued)**

Action	Direction	Description

Account	Ingress / Egress	counts traffic matching the specified flows:
		• spawn a dynamic policer per direction (if not already present)
		• if no rate-limit action is specified in the PCC rule, set the PIR/CIR value of the dynamic policer to **max**
		• map the flows to the policer
		The forwarded octets and packets statistics of the dynamic policer are included in subscriber service accounting.
		CLI equivalent:
		`config>qos`
		`  sap-ingress | sap-egress <id> create`
		`    policer 1  # dynamic policer`
		`      rate max cir max`
		`    exit`
		`    ip-criteria | ipv6-criteria`
		`      entry <id> create`
		`        match`
		`          <5-tuple | dscp>`
		`        exit`
		`        action policer 1`
		`      exit`
		`    exit`

Table 92 **Subscriber Service PCC Rule Actions Resulting in QoS Policy Changes (Continued)**

Action	Direction	Description

Forward	Ingress/Egress	Creates an entry in the QoS policy to forward the traffic without explicit QoS action. Matching traffic will not match on the next entry (match and exit behavior). In case of overlapping flows, such as account all traffic except flow 1 and flow 2, the more specific flows must be explicitly forwarded.
		CLI equivalent:
		`config>qos`
		`  sap-ingress | sap-egress <id> create`
		`    ip-criteria | ipv6-criteria`
		`      entry <id> create`
		`        match`
		`          <5-tuple | dscp>`
		`        exit`
		`        action`
		`      exit`
		`    exit`

Table 93 **Subscriber Service PCC Rule Actions Resulting in Filter Changes**

Action	Direction	Description

Forward/Drop	Ingress/Egress	creates a filter entry to forward or drop the traffic
		CLI equivalent:
		`config>filter`
		`  ip-filter | ipv6-filter <id> create`
		`      entry <id> create`
		`        match`
		`          <5-tuple | dscp>`
		`        exit`
		`        action`
		`          forward | drop`
		`        exit`
		`      exit`
		`    exit`

Table 93 **Subscriber Service PCC Rule Actions Resulting in Filter Changes (Continued)**

Action	Direction	Description

Redirect to an IP next-hop	Ingress	redirect the traffic to the specified IP next-hop
		CLI equivalent:
		`config>filter`
		`  ip-filter | ipv6-filter <id> create`
		`  entry <id> create`
		`  match`
		`  <5-tuple | dscp>`
		`  exit`
		`  action`
		`  forward next-hop <ip-address>`
		`  exit`
		`  exit`
		`  exit`

Redirect to a routing instance	Ingress	redirect the traffic to the specified routing instance
		CLI equivalent:
		`config>filter`
		`  ip-filter | ipv6-filter <id> create`
		`  entry <id> create`
		`  match`
		`  <5-tuple | dscp>`
		`  exit`
		`  action`
		`  forward router <router-instance>`
		`  exit`
		`  exit`
		`  exit`

Table 93 **Subscriber Service PCC Rule Actions Resulting in Filter Changes (Continued)**

Action	Direction	Description

HTTP redirect	Ingress	HTTP redirection to the specified URL
		CLI equivalent:
		`config>filter`
		`  ip-filter | ipv6-filter <id> create`
		`  entry <id> create`
		`  match`
		`  <5-tuple | dscp>`
		`  exit`
		`  action`
		`  http-redirect <rdr-url-string>`
		`  exit`
		`  exit`
		`  exit`

Figure 158 and Figure 159 show the actions that can be combined in a single ingress or egress PCC rule.

Figure 158 **Supported Combinations of Ingress PCC Rule Actions**

INGRESS supported combinations of PCC rule actions		QoS				Filter				Gate
		Forward	Rate limit	FC change	Account / UM	Forward	Drop	redirect next-hop or router	redirect URL	Flow Status (1)
QoS	Forward	✓	✓ (3)	✓ (3)	✓ (3)	✓	–	✓	–	✓
	Rate limit	✓ (3)	✓	✓	✓ (4)	✓	–	✓	–	✓
	FC change	✓ (3)	✓	✓	✓	✓	–	✓	–	✓
	Account / UM	✓ (3)	✓ (4)	✓	✓	✓	–	✓	–	✓
Filter	Forward	✓	✓	✓	✓	✓	–	✓ (2)	✓ (2)	✓
	Drop	–	–	–	–	–	✓	–	–	✓
	redirect next-hop or router	✓	✓	✓	✓	✓ (2)	–	✓	–	✓
	redirect URL	–	–	–	–	✓ (2)	–	–	✓	✓
Gate	Flow Status (1)	✓	✓	✓	✓	✓	✓	✓	✓	✓

0961

Figure 159 **Supported Combinations of Egress PCC Rule Actions**

EGRESS supported combinations of PCC rule actions		QoS				Filter		Gate
		Forward	Rate limit	FC change	Account / UM	Forward	Drop	Flow Status (1)
QoS	Forward	✓	✓(3)	✓(3)	✓(3)	✓	-	✓
	Rate limit	✓(3)	✓	✓	✓(4)	✓	-	✓
	FC change	✓(3)	✓	✓	✓	✓	-	✓
	Account / UM	✓(3)	✓(4)	✓	✓	✓	-	✓
Filter	Forward	✓	✓	✓	✓	✓	-	✓
	Drop	-	-	-	-	-	✓	✓
Gate	Flow Status (1)	✓	✓	✓	✓	✓	✓	✓

0962

Notes:

1. The Flow Status can only be set via Gx.

   For RADIUS subscriber service-based PCC rules, the Flow Status is fixed to Enabled.
2. The filter action forward is ignored (not installed) when combined with the following filter actions: redirect next-hop, redirect router, or http-redirect.
3. The QoS action forward is ignored (not installed) when combined with the following QoS actions: rate-limit, FC change, account/Usage Monitoring (UM).
4. When the QoS action rate-limit and QoS action account or Usage Monitoring (UM) are combined, only a single dynamic policer is installed, which is used for both rate-limiting and obtaining forward statistics for accounting or usage monitoring.

9.18.6.2 PCC Rule Instantiation

A PCC rule can result in one or more IPv4/IPv6 filter and/or QoS policy IPv4/IPv6 criteria entries. This is transparent to the operator.

- A PCC rule is split into IPv4 filter entries, IPv6 filter entries, SAP ingress QoS IP or IPv6 criteria, and SAP egress QoS IP or IPv6 criteria.
- Each entry is inserted into the corresponding policy at a reserved range for dynamic PCC rule inserts. Within the reserved range, the (optional) precedence value for the rule is taken into account for the relative order of different PCC rules.
- The QoS rate-limit and account actions will spawn a dynamic policer from a reserved range in the QoS policy. A template configuration provides dynamic policer parameters such as hierarchical policing parent, burst sizes (MBS, CBS), statistics mode and packet byte offset. A maximum of one dynamic policer is instantiated per PCC rule. There is a maximum of 63 dynamic policers per direction and per SLA profile instance. The output queue for PCC rule traffic mapped in the dynamic policer is determined by a mechanism called forwarding class inheritance: the output queue is the same queue that would be used if a packet with the same forwarding class as the PCC rule packet was classified using the applied QoS policy. The resulting output queue can be a local subscriber queue (when the FC is mapped to a queue or when the FC is mapped to a policer at egress and the policer is mapped to a local queue), a shared queue (when the FC is mapped to a policer at ingress) or a queue-group queue (when the FC is mapped to a policer at egress).
- Optimal policy and rule sharing is achieved via QoS and filter policy cloning and internal PCC rule optimizations. The mechanisms are the same as for Gx-initiated PCC rules as described the Generic Policy Sharing and Rule Sharing and Gx Rule Ordering.

  PCC rule sharing can only happen when the content is the same: identical name, direction, precedence value, set of flows, and set of actions. PCC rules with the same content will have the same PCC rule ID.

  Filter and QoS policy clones that result from PCC rule instantiation can be recognized by a filter ID or QoS policy ID in the format 1 to 65535:P1 to 4096; for example, filter ip 10:p3
- PCC rules can be inactive if the corresponding host type is not present. For example, a PCC rule-based subscriber service with an IPv6 filter action can be activated on an IPoE session while there is no IPv6 host instantiated on the session. When the IPv6 host is later created, the PCC rule is activated.

The following initial configuration is required prior to activating PCC rule-based subscriber services.

- To install PCC rule QoS actions, a non-default ingress and egress QoS policy with a **sub-insert-shared-pccrule** range configured must be associated with the IPoE or PPPoE session (the default QoS policy cannot be modified). If a rate-limit or account action is needed, a dynamic policer range must also be configured. Additional dynamic policer parameters are optional.

```
configure qos
```

```
sap-ingress <policy-id> create
    sub-insert-shared-pccrule start-entry <entry-id> count <count>
    dynamic-policer
        range start-entry <policer-id> count <count>
        packet-byte-offset {add <add-bytes> | subtract <sub-bytes>}
        mbs <size> [bytes|kilobytes]
        cbs <size> [bytes|kilobytes]
        parent <arbiter-name> [weight <weight-level>] [level <level>]
        stat-mode <stat-mode>
    exit
exit
```

- To install PCC rule filter actions, an IPv4 and/or IPv6 filter with **sub-insert-shared-pccrule** range configured must be associated with the IPoE or PPPoE session.

Although a PCC rule-based subscriber service is activated on a PPPoE or IPoE session, the actions are applied at the SLA profile instance and/or subscriber host level.

- PCC rule QoS actions result in QoS policy clones that are applied at the SLA profile instance level. Traffic from all subscriber hosts and sessions sharing the SLA profile instance will be subject to the specified actions.
- PCC rule filter actions result in IPv4 or IPv6 filter clones that are applied at the subscriber host level. Only traffic from the subscriber host of the same type (IPv4 or IPv6) that belongs to the PPPoE or IPoE session will be subject to the specified actions.

A PCC rule with flow match criteria that are not explicitly IPv4 or IPv6 results in both IPv4 and IPv6 match criteria being installed; for example, destination address = any.

Filter actions are executed before QoS actions. If an IP flow is rate-limited, it should pass the IPv4 or IPv6 filter first. Adding a QoS action rate limit to a PCC rule will not automatically insert a corresponding forward entry in an IP filter. When needed, this must be done explicitly by the operator with a filter forward action. For example, an IP filter with the default action drop and a number of explicit forward entries is applied to an IPoE session. A new IP flow must be rate-limited and accounted for. The PCC rule should include match criteria for the IP flow and a QoS action rate limit, QoS action account, and filter action forward. Without the filter action forward, the IP flow would be dropped by the default action in the filter policy.

```
config>filter
    ip-filter <filter-id> create
        sub-insert-shared-pccrule start-entry <entry-id> count <count>
    exit
    ipv6-filter <filter-id> create
        sub-insert-shared-pccrule start-entry <entry-id> count <count>
    exit
```

9.18.6.3 PCC Rules in a Subscriber Service

The example below shows a pseudo-language representation of PCC rules in a subscriber service.

- A subscriber service can contain multiple PCC rules. Because a PCC rule is unidirectional, including an ingress and an egress PCC rule enables subscriber service accounting of bidirectional flows.
- A PCC rule can contain multiple flows. Flows in a PCC rule can be a mix of IPv4 and IPv6 flows.

```
subscriber-service {
    name = <subsvc name>
    operation = add | delete
    acct-stats-type = off | volume-time | time
    acct-interval = <value>
    type = <string>
    type-conflict-action = keep-old | keep-new | none
    pcc-rule {
        name = <name>
        direction = ingress | egress
        flow = <5-tuple> | <dscp>
        flow = <5-tuple> | <dscp>
           …
        flow = <5-tuple> | <dscp>
        action = <action>
        action = <action>
           …
        action = <action>
        precedence = <value>
    }
    …
    pcc-rule {
        name = <name>
        direction = ingress | egress
        flow = <5-tuple> | <dscp>
        flow = <5-tuple> | <dscp>
           …
        flow = <5-tuple> | <dscp>
        action = <action>
        action = <action>
           …
        action = <action>
        precedence = <value>
    }
}
```

The sub_svc.pccrule TLV in the subscriber services Python script adds a PCC rule to the subscriber service, as shown in the output example below:

Actual values are used to populate the subscriber service data structure in this example; typically, these values are sent as parameters in subscriber service-specific VSAs.

```
# Python example - PCC rules subscriber service
service = []
sub_svc.add_to_service(service, sub_svc.operation, sub_svc.operation_add)
sub_svc.add_to_service(service, sub_svc.name, 'subsvc-pccrule-1')
sub_svc.add_to_service(service, sub_svc.type, 'pccrule')
sub_svc.add_to_service(service, sub_svc.type_conflict_action, sub_svc.type_conflict_
action_none)
flow_i = []
sub_svc.flow.add_to_flow(flow_i, sub_svc.flow.dst_ip, '10.1.1.0/24')
rule_i = []
sub_svc.pccrule.add_to_pccrule(rule_i, sub_svc.pccrule.name, 'pcc-rule-1-i')
sub_svc.pccrule.add_to_pccrule(rule_i, sub_svc.pccrule.precedence, 10)
sub_svc.pccrule.add_to_pccrule(rule_i, sub_svc.pccrule.direction, sub_svc.pccrule.di
rection_ingress)
sub_svc.pccrule.add_to_pccrule(rule_i, sub_svc.pccrule.flow, flow_i)
sub_svc.pccrule.add_to_pccrule(rule_i, sub_svc.pccrule.qos_action_rate_limit_pir, 10
00)
sub_svc.pccrule.add_to_pccrule(rule_i, sub_svc.pccrule.qos_action_account, True)
flow_e = []
sub_svc.flow.add_to_flow(flow_e, sub_svc.flow.src_ip, '10.1.1.0/24')
rule_e = []
sub_svc.pccrule.add_to_pccrule(rule_e, sub_svc.pccrule.name, 'pcc-rule-1-e')
sub_svc.pccrule.add_to_pccrule(rule_e, sub_svc.pccrule.precedence, 10)
sub_svc.pccrule.add_to_pccrule(rule_e, sub_svc.pccrule.direction, sub_svc.pccrule.di
rection_egress)
sub_svc.pccrule.add_to_pccrule(rule_e, sub_svc.pccrule.flow, flow_e)
sub_svc.pccrule.add_to_pccrule(rule_e, sub_svc.pccrule.qos_action_rate_limit_pir, 50
00)
sub_svc.pccrule.add_to_pccrule(rule_i, sub_svc.pccrule.qos_action_account, True)
sub_svc.add_to_service(service, sub_svc.pccrule, rule_i)
sub_svc.add_to_service(service, sub_svc.pccrule, rule_e)
sub_svc.add_to_service(service, sub_svc.acct_stats_type, 2)
sub_svc.add_to_service(service, sub_svc.acct_interval, 3600)
sub_svc.commit_service(service)
```

In the above subscriber service example, two PCC rules are installed, each with one flow:

- pcc-rule-1-i: ingress traffic to destination 10.1.1.0/24 is rate-limited to 1 Mb/s
- pcc-rule-1-e: egress traffic from 10.1.1.0/24 is rate-limited to 5 Mb/s

Subscriber service instance volume-time accounting is enabled. The volume counters include the forwarded octets and packets of the dynamic policers installed for the above rules and count the traffic matching the flows.

The counters reported in PCC rule-based subscriber services RADIUS accounting are determined by the PCC rule QoS account action.

- If at least one PCC rule in the subscriber service has the QoS account action enabled (pccrule.qos_action_account = True), then the volume counters contain the sum of the dynamic policer forwarded octets and packets statistics of all the PCC rules in the subscriber service with pccrule.qos_action_account = True.

- If all PCC rules in the subscriber service have the QoS account action disabled (pccrule.qos_action_account = False), then the volume counters contain the aggregated forwarded octets and packets of the parent PPPoE or IPoE session SLA profile instance since the start of the subscriber service.

PCC rule-based subscriber services are not stored in the subscriber-mgmt persistency file.

9.18.6.4 Interaction of the PPPoE or IPoE Session QoS Model and PCC Rule-based Subscriber Services

PCC rule-based subscriber services with QoS actions interact with the classification and QoS forwarding mechanisms. This section describes how this affects the parent RADIUS accounting volume counters.

For subscriber service PCC rule QoS actions that do not result in the instantiation of a dynamic policer (such as a change of forwarding class or forward), the PCC rule matched traffic is included in the parent accounting session volume counters. This is shown in Figure 160, where the forwarding class is mapped to a subscriber queue, and in Figure 161 where the forwarding class is mapped to a subscriber policer.

Figure 160 **PCC Rule-based Subscriber Service—QoS Interaction: No Dynamic Policer, FC → Queue**

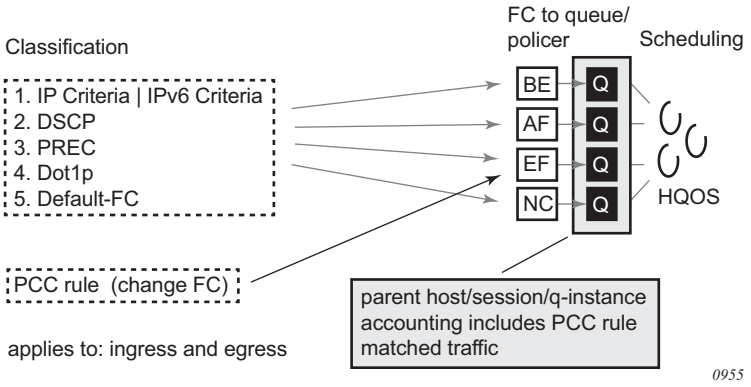

Figure 161 **PCC Rule-based Subscriber Service—QoS Interaction: No Dynamic Policer, FC → Policer → Queue-group**

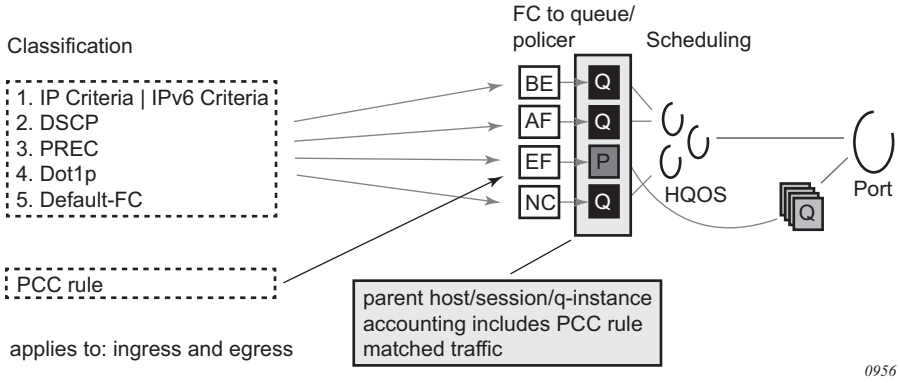

For subscriber service PCC rule QoS actions that result in the instantiation of a dynamic policer (such as rate-limit or account), the dynamic policer counters are not included in the aggregate counters nor are they reported as separate detailed policer statistics. Instead, the traffic matching the PCC rules is counted in the output queues that correspond to the forwarding class of the packets.

- On ingress, the dynamic policer PCC rule traffic is never included in the parent host, session, or queue instance accounting session counters. Ingress policed traffic always uses the ingress shared policer output queues, as shown in Figure 162.

Figure 162 **PCC Rule-based Subscriber Service—QoS Interaction: Dynamic Policer → Ingress Shared Policer Output Queues**

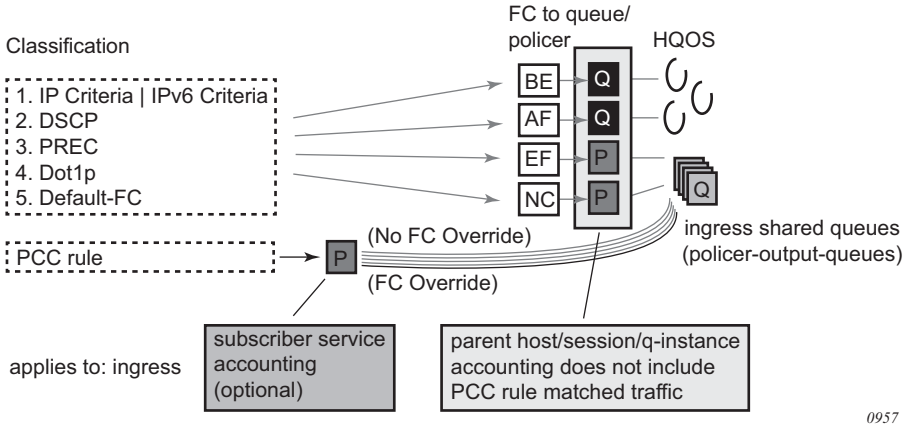

- To include the egress dynamic policer PCC rule traffic in the parent host, session, or queue instance accounting session counters, the dynamic policer must use a local subscriber output queue, as shown in Figure 163.

Figure 163 **PCC Rule-based Subscriber Service—QoS Interaction: Dynamic Policer, FC → Queue (Egress)**

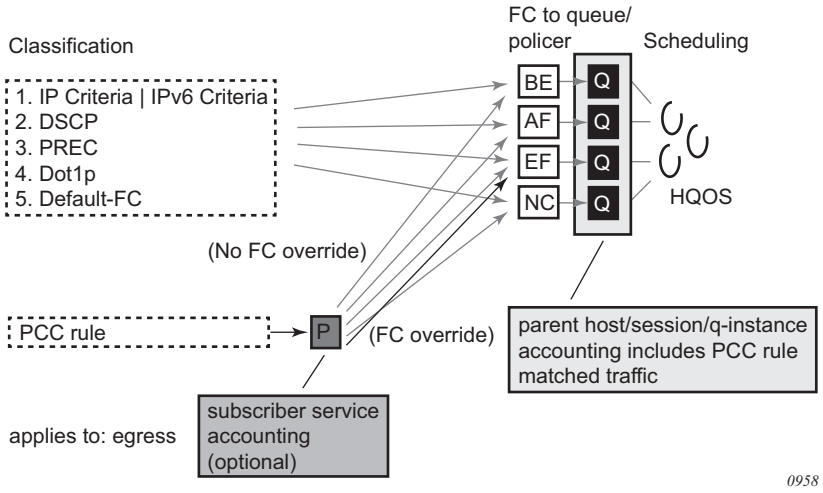

- To exclude the egress dynamic policer PCC rule traffic from the parent host, session, or queue instance accounting session counters, the dynamic policer must use a queue-group output queue, as shown in Figure 164. The dynamic policer traffic inherits the policer-to-output queue mapping from the static policer that corresponds to the forwarding class of the packet. The following SAP egress configuration example uses the default policer output queue-group:

```
config>qos
    sap-egress 10 create
       sub-insert-shared-pccrule start-entry 200 count 10
        dynamic-policer
            range start-entry 10 count 10
        exit
        policer 1 create
        exit
        fc be create
            policer 1
        exit
    exit
```

If a packet is classified as FC = Best Effort (BE) and matches a PCC rule with rate-limit action only (no FC change), the traffic will hit the PCC rule dynamic policer and then the queue-group queue associated with policer 1 (the static policer for FC = BE).

Figure 164 **PCC Rule-based Subscriber Service—QoS Interaction: Dynamic Policer, FC → Policer → Policer Output Queues (Egress)**

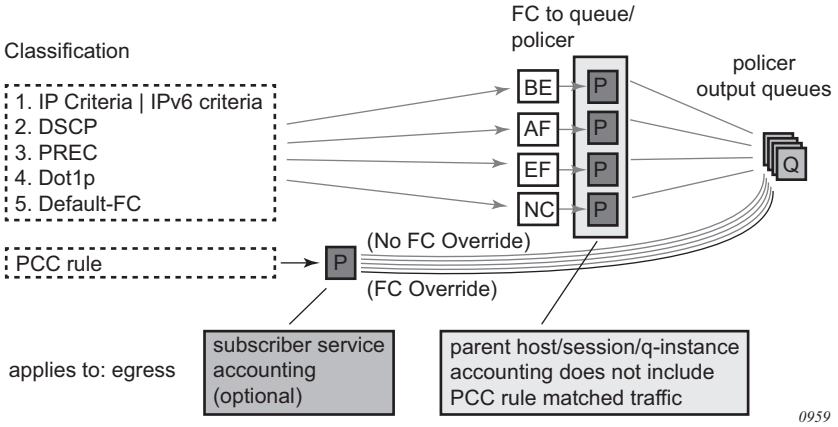

- A special case occurs when, at egress, the forwarding class maps to a static policer and then to a local subscriber queue. Traffic for that forwarding class hitting a dynamic policer will use the local subscriber queue as the output queue. In this case, the dynamic policer PCC rule traffic is included in the parent host, sessions or queue instance accounting session counters. This is shown in Figure 165.

Figure 165 **PCC Rule-based Subscriber Service—QoS Interaction: Dynamic Policer, FC → Policer → Local Queue (Egress)**

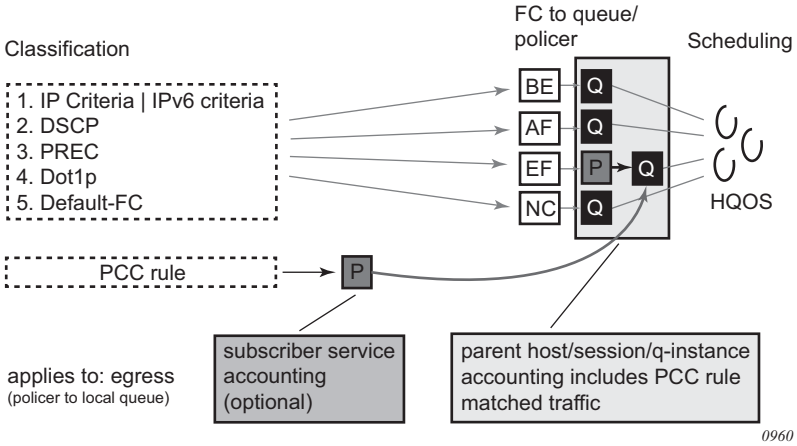

9.18.6.5 PCC Rules on HSMDAv2

This section describes the differences in PCC rule-based subscriber services when applied to subscriber sessions terminating on an HSMDAv2 card, further referred to as HSMDA.

The activation of PCC rule-based subscriber services is explained in the previous sections.

The filter policy cloning resulting from the instantiation of PCC rule filter actions is the same on the HSMDA:

- IP or IPv6 filter policies applied at the sla-profile ingress or egress contexts can be cloned when PCC rules with filter actions are instantiated on HSMDA subscriber sessions.

The QoS policy cloning resulting from the instantiation of PCC rule QoS actions is different on HSMDAs because of the specific egress QoS architecture:

- **ingress** — The QoS policy applied at the sub-profile contains the active configuration and can be cloned when ingress PCC rules with QoS actions are instantiated on HSMDA subscriber sessions. When a PCC rule clone exists, the QoS policy applied at the sub-profile ingress cannot be changed.
- **egress**: The QoS policy applied at the sla-profile contains the active configuration except for the HSMDA queues. This QoS policy can be cloned when egress PCC rules with QoS actions are instantiated on HSMDA subscriber sessions. When a PCC rule clone exists, the QoS policy associated at the sla-profile egress cannot be changed.

Table 94 summarizes the configuration of HSMDA QoS and filter policies:

Table 94 **HSMDA QoS and Filter Policy Configurations**

Direction	Profile
SLA Profile	
Ingress	CLI context: **config>subscr-mgmt>sla-prof>ingress** • ip-filter • ipv6-filter • qos <policy-id> (This configuration is ignored for HSMDA.)

Table 94 **HSMDA QoS and Filter Policy Configurations (Continued)**

Direction	Profile
Egress	CLI context: **config>subscr-mgmt>sla-prof>ingress** • ip-filter • ipv6-filter • qos <policy-id> – Egress reclassification: precedence, DSCP – Forwarding Class to HSMDA queue mapping – ip-criteria and ipv6-criteria – dynamic policer configuration – sub-insert-shared-pccrule – no active policer configuration allowed: • Forwarding Class to policer mapping • ip-criteria and ipv6-criteria action policer <id> – The IOM queue configuration is ignored
Sub Profile	
Ingress	CLI context: **config>subscr-mgmt>sub-prof>hsmda>ingress-qos** • qos <policy-id> – IOM queue and policer configuration – classification: dot1p, precedence, DSCP, default-fc, default-priority – Forwarding class to queue/policer mapping – ip-criteria and ipv6-criteria, including action policer <id> – dynamic policer configuration – sub-insert-shared-pccrule

Table 94 **HSMDA QoS and Filter Policy Configurations (Continued)**

Direction	Profile
Egress	CLI context: **config>subscr-mgmt>sub-prof>hsmda>egress-qos** • qos <policy-id> – No active policer configuration allowed: • Forwarding Class to policer mapping • ip-criteria and ipv6-criteria action policer <id> • HSMDA queue configuration • Other configurations are ignored

The following restrictions apply for PCC rule based subscriber services on HSMDA subscriber sessions:

- An egress rate limit (PIR/CIR) action is not supported. Egress dynamic policers on HSMDA are always installed with PIR equaling max and CIR equaling 0. The CBS value for dynamic policers is ignored and it is recommended to keep the default MBS value.
- There is no support for hierarchical policing on egress dynamic policers. The subscriber host or session setup fails when a sub-profile is assigned containing an egress policer-control-policy configuration.
- Only a single SLA Profile Instance (SPI) can be active for a subscriber when a PCC-rule subscriber service containing egress QoS actions is active or activated:
  - The activation of a PCC rule-based subscriber service containing egress QoS actions fails when more than one SPI is active for the subscriber.
  - A new subscriber host or session setup fails when a new SPI is required while a PCC rule based subscriber service containing egress QoS actions is active.
  - A PCC rule based subscriber service containing egress QoS actions must be removed before the SLA profile of a subscriber can be changed.

9.18.6.5.1 Interaction between PCC Rule-Based Subscriber Services and RADIUS Queue-Instance, Host or Session Accounting Volume Counters

PCC rule based subscriber services with QoS actions interact with the classification and QoS forwarding mechanisms. This section describes how this affects the parent RADIUS accounting volume counters.

For subscriber service PCC rule QoS actions that result in the instantiation of a dynamic policer (such as rate-limit or account), the dynamic policer counters are not included in the aggregate counters nor are they reported as separate detailed policer stats. Instead the traffic matching the PCC rules is counted in the output queues that correspond with the Forwarding Class of the packets:

- On ingress, the dynamic policer PCC rule traffic is never included in the parent host, session or queue-instance accounting session counters. Ingress policed traffic always uses the ingress shared policer output queues. This is the same as for non-HSMDA scenarios. See Figure 162.
- On egress, the dynamic policer PCC rule traffic is always included in the parent host, session or queue-instance accounting session counters. Egress policed traffic always uses the egress HSMDA queues as shown in Figure 166.

Figure 166 **PCC Rule-Based Subscriber Service – HSMDA Egress Accounting**

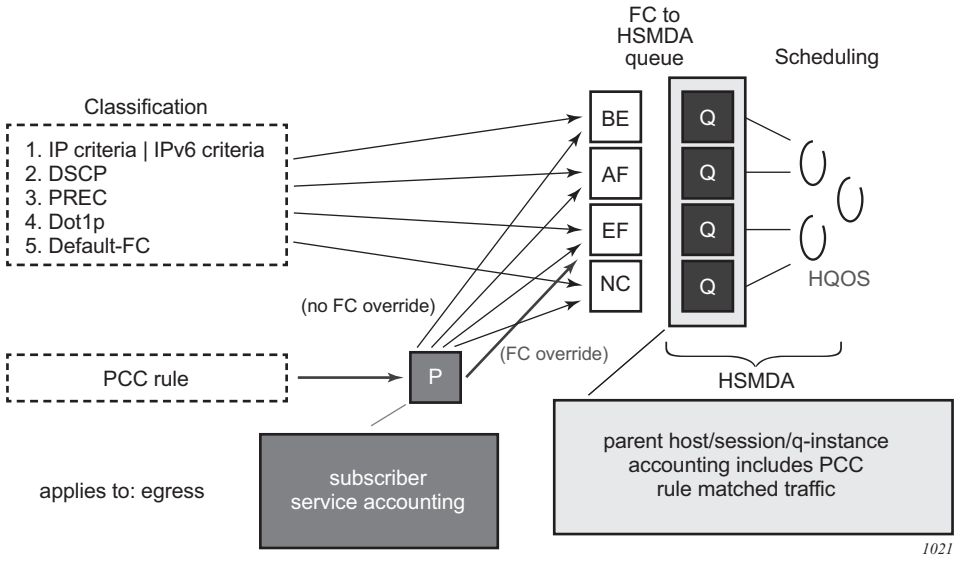

9.18.6.6 PCC Rule-based Subscriber Service Activation Failures

PCC rule-based subscriber service activation failures can be categorized as failures detected in the subscriber services Python script as runtime errors (see Table 95) and failures detected in the Enhanced Subscriber Management (ESM) processing (see Table 96).

Table 95 **PCC Rule-based Subscriber Service: Python Runtime Errors**

Python Runtime Errors
Pcc rule name type must be a string
Pcc rule name value too long
Pcc rule qos action account value must be True(1) or False(0)
Pcc rule filter action redirect to nexthop_v4/v6 type must be a string
Pcc rule filter action redirect to nexthop_v4/v6 value must be a valid IPv4/v6 address
Pcc rule filter action http redirect value is not a valid URL
Pcc rule filter action http redirect type must be a string
Pcc rule filter action http redirect value too long (> 255)
Pcc rule qos action fwd class change value is invalid
Pcc rule qos action rate limit type must be an int
Pcc rule precedence type must be an int
Pcc rule flow dscp match type must be a string
Pcc rule flow dscp match value is invalid
Pcc rule flow src/dst_ip match type must be a string: <ipv4-address>|<ipv6-address>|any
Pcc rule flow src/dst_ip match value is not a valid IP address: <ipv4-address>|<ipv6-address>|any
Pcc rule flow port match type must be a string: <port>[-<port>] with port [0..65535]
Pcc rule flow protocol match type must be an int
Pcc rule flow protocol match value must be less than 255
Pcc rule must have a name

Table 96 **PCC Rule-based Subscriber Service: ESM Failures**

ESM Failures
ESM —RADIUS Decoding Failures
The PCC rule precedence must be in the range 0 to 65535.
The PCC rule name has a maximum of 100 displayable characters.
The PCC rule redirect URL has a maximum of 255 displayable characters.
All flows in a PCC rule must have the same direction (ingress or egress).
A PCC rule has a maximum of 128 flows.
A PCC rule name can occur only once in a RADIUS message.
A flow in a PCC rule cannot have a mix of IPv4 and IPv6 addresses (for example **src-ip** and **dst-ip**).
ESM—Processing Failures
If a PCC rule contains a direction-specific action (such as a redirect), it must contain at least one flow in that direction.
If a PCC rule contains only IPv4 actions (such as a redirect to an IPv4 next hop), it must contain at least one IPv4 flow. This also applies to IPv6.
The combinations of PCC rule actions must be supported (see Figure 158 and Figure 159).
There must be at least one flow and at least one action per PCC rule.
There is a maximum of 64 PCC rules per host or session.
There are not enough filter or QoS resources to create policy clones or apply them to the host or session.
The filter or QoS policy clone cannot be created (for example, the redirect service does not exist).
Simultaneous provisioning of PCC rules from Gx and RADIUS is operationally blocked per subscriber session/host.

9.18.7 Combined Subscriber Services

A subscriber service instance can be built with a combination of QoS override actions and PCC rules. In this case, the reported volume counters for subscriber service statistics type volume-time are determined by the PCC rule account action, as shown in Table 97.

Table 97 **Subscriber Service Accounting—Reported Volume Counters**

QoS Override	PCC Rule	Reported Volume Counters
No	No	Aggregated SLA profile queue and policer statistics
Yes	No	Aggregated SLA profile queue and policer statistics
No	Yes, “pccrule.qos_action_account = False” for all PCC rules	Aggregated SLA profile queue and policer statistics
No	Yes, “pccrule.qos_action_account = True” for at least one PCC rule	Dynamic policer statistics for PCC rules with “pccrule.qos_action_account = True”
Yes	Yes, “pccrule.qos_action_account = False” for all PCC rules	Aggregated SLA profile queue and policer statistics
Yes	Yes, “pccrule.qos_action_account = True” for at least one PCC rule	Dynamic policer statistics for PCC rules with “pccrule.qos_action_account = True”

9.18.8 Subscriber Services Python API

The SR OS alc.sub_svc Python module offers functions and data structures to describe, activate, and deactivate a subscriber service.

9.18.8.1 Common Subscriber Services Python API

Table 98 lists the subscriber service functions in the alc.sub_svc module.

Table 98 **Subscriber Service Functions in the alc.sub_svc Module**

sub_svc Functions	Description
sub_svc.add_to_service (svc, sub_svc TLV, value)	appends a TLV to the service list. The service list describes the subscriber service and should be passed to the sub_svc.commit_service() to activate or deactivate the subscriber service. Parameters: svc (type = list): service list that describes the subscriber service. sub_svc TLVs are appended to this list with the sub_svc.add_to_service() function. sub_svc TLV (type int): TLV that will be appended to the service list value (type as defined for the sub_svc TLV): the value of the sub_svc TLV that will be appended to the service list
sub_svc.commit_service (svc)	creates the required internal VSAs based on the TLVs provided in the svc list. Parameters: svc (type = list): service list that describes the subscriber service. sub_svc TLVs are appended to this list with the sub_svc.add_to_service() function. The service list should be passed to sub_svc.commit_service() to activate or deactivate the subscriber service.

Table 99 lists the subscriber service TLVs in the alc.sub_svc module.

Table 99 **Subscriber Service TLVs in the alc.sub_svc Module**

sub_svc TLV	Activate/ Deactivate	TLV Details	
name (string)	M/M	Purpose	The unique subscriber service instance identifier (key). This field will be matched for a subscriber service deactivate request. This field could, for example, be populated with the service-name and corresponding parameter list. This value is also echoed in the Alc-Sub-Serv-Session attribute in accounting messages. When not specified, the subscriber service activation fails and an event log is generated: WARNING: SVCMGR #2511 Base RADIUS CoA Error “Problem encountered in Subscriber Management, while processing a CoA request from a RADIUS server: Could not decode RADIUS Attribute “Sub Service””.
		Value	Free format string (max. length = 1000 bytes)
		Default	Empty string
operation (int)	M/O	Purpose	Specifies if the referenced subscriber service should be activated or deactivated
		Value	operation_add (1): activate the subscriber service instance operation_del (2): deactivate the subscriber service instance
		Default	operation_del
acct_stats_type (int)	O/n.a.	Purpose	Defines if RADIUS accounting should be enabled for this subscriber service instance. If enabled, the accounting mode (Time or Volume and Time) is specified. Values as defined for Alc-Sub-Serv-Acct-Interim-Ivl VSA
		Value	Off (1) Volume-time (2) Time (3)
		Default	Off

Table 99 **Subscriber Service TLVs in the alc.sub_svc Module (Continued)**

sub_svc TLV	Activate/ Deactivate	TLV Details	
acct_interval (int)	O/n.a.	Purpose	Defines the RADIUS interim accounting update interval for this subscriber service instance Values as defined for Alc-Sub-Serv-Acct-Interim-Ivl VSA
		Value	0 (no interim updates) 1 to 299 (rounded to a maximum 300 seconds) 300 to 15552000 (override the local configured update-interval for this subscriber service instance)
		Default	Update interval from the parent subscriber RADIUS accounting policy
type (string)	O/n.a.	Purpose	Grouping of subscriber service instances that belong to the same PPPoE or IPoE session
		Value	Free format string (max. length = 255 bytes)
		Default	empty string
type_conflict_action (int)	O/n.a.	Purpose	Defines the action when another subscriber service instance of the same type is already activated for the same PPPoE or IPoE session.
		Value	type_conflict_action_keep_old (1): reject the new subscriber service instance type_conflict_action_keep_new (2): deactivate the old and activate the new subscriber service instance type_conflict_action_none (3): allow multiple subscriber service instances of this type
		Default	type_conflict_action_none

Note: M=Mandatory, O=Optional, n.a.=ignored

9.18.8.2 Subscriber Service QoS Override Python API

Table 100 lists the QoS override TLVs in the alc.sub_svc module.

Table 100 **QoS Override TLVs in the alc.sub_svc Module**

sub_svc TLV	Activate/ Deactivate	TLV Details	
qos_override (string)	O/n.a.	Purpose	Adds a QoS override to the subscriber service. Multiple qos_override TLVs can be added in a single subscriber service instance.
		Value	As defined for the RADIUS Alc-Subscriber-QoS-Override VSA [26-6527-126]. Refer to the *7750 SR and VSR RADIUS Attributes Reference Guide* for details.
		Default	Not included

Note: M=Mandatory, O=Optional, n.a.=ignored

9.18.8.3 Subscriber Service PCC Rules Python API

Table 102 lists the PCC rule PLVs in the alc.sub_svc module.

Table 101 **PCC Rule TLVs in the alc.sub_svc Module**

sub_svc TLV	Activate/ Deactivate	TLV Details	
pccrule (list)	O/n.a.	Purpose	Adds a PCC rule to the subscriber service. Multiple PCC rule TLVs can be added in a single subscriber service instance.
		Value	A PCC rule list describing the PCC rule with PCC rule TLVs such as name, precedence, direction, flows, and actions PCC rule TLVs are appended to the PCC rule list with the sub_svc.pccrule.add_to_pccrule() function.
		Default	Not included

Note: M=Mandatory, O=Optional, n.a.=ignored

Table 102 lists the alc.sub_svc.pccrule function.

Table 102 **alc.sub_svc.pccrule Function**

alc.sub_svc.pccrule Function	Description
sub_svc.pccrule.add_to_pccrule (pccrule, pccrule TLV, value)	Appends a PCC rule TLV such as name, precedence, flow, or action to the PCC rule list. The PCC rule list describes the PCC rule and can be added to a subscriber service with the sub_svc.add_to_service() function. Parameters: pccrule (type = list): PCC rule list that describes the PCC rule. PCC rule TLVs are appended to this list with the sub_svc.pccrule.add_to_pccrule() function. pccrule TLV (type int): PCC rule TLV that will be appended to the PCC rule list value (type as defined for the PCC rule TLV): the value of the PCC rule TLV that will be appended to the PCC rule list

Table 103 lists the PCC rules TLVs in the alc.sub_svc.pccrule module.

Table 103 **PCC Rule TLVs in the alc.sub_svc.pccrule Module**

PCC Rule TLV	M|O	TLV Details	
pccrule.name (String)	M	Purpose	Specifies the name of the PCC rule A PCC rule with the same name and same or different content can only be applied once on a single parent PPPoE or IPoE session. A PCC rule with the same name and same or different content can be applied on different parent PPPoE or IPoE sessions. Rules with the same name but different content will get a different PCC rule identifier (rule id).
		Value	Free format string (maximum length = 100 bytes)
		Default	n/a

Table 103 **PCC Rule TLVs in the alc.sub_svc.pccrule Module (Continued)**

PCC Rule TLV	M|O	TLV Details	
pccrule.precedence (Integer)	O	Purpose	Specifies the precedence value for the PCC rule. The precedence defines a relative order of the different PCC rules: a rule with a lower precedence value will be applied before a rule with a higher precedence value. Rules with the same precedence and rules without precedence can be automatically optimized; the relative order in which they are applied is determined by the system for optimal sharing.
		Value	0 to 65535
		Default	n/a These rules are applied at the end.
pccrule.direction (Integer)	M	Purpose	Specifies the direction of the PCC rule: ingress or egress
		Value	direction_ingress (1)
			direction_egress (2)
		Default	n/a
pccrule.flow (list)	M	Purpose	Adds a flow to the PCC rule. At least one flow must be added to a PCC rule. Multiple flow TLVs can be added to a PCC rule.
		Value	A flow list describing the flow with flow TLVs such as dscp, protocol, src-ip, dst-ip, src-port, and dts-port Flow TLVs are appended to the flow list with the sub_svc.flow.add_to_flow() function.
		Default	Not included

Table 103 **PCC Rule TLVs in the alc.sub_svc.pccrule Module (Continued)**

PCC Rule TLV	M|O	TLV Details	

pccrule.qos_action_ account (Boolean)	O (1)	Purpose	PCC rule action: account
			Can be applied on ingress and egress
			Results in IPv4 and/or IPv6 criteria entry in QoS policies.
			When set to True:
			if no rate-limit action is specified, a dynamic policer with pir=cir=max is instantiated for all flows in the PCC rule
			CLI equivalent:
			`policer 1 # dynamic policer`
			`rate max cir max`
			`exit`
			`entry 10 create`
			`match`
			`…`
			`exit`
			`action policer 1`
			`exit`
			The forwarded octets and packets statistics of the dynamic policer associated with this PCC rule are included in subscriber service accounting.

		Value	True (1) False (0)
		Default	False

pccrule.qos_action_ change_fc (string)	O (1)	Purpose	PCC rule action: change the forwarding class
			Can be applied on ingress and egress.
			Results in IPv4 and/or IPv6 criteria entry in QoS policies
			CLI equivalent:
			`entry 10 create`
			`match`
			`…`
			`exit`
			`action fc <fc-name>`
			`exit`

		Value	String with fixed format forwarding class name: “be”, “l2”, “af”, “l1”, “h2”, “ef”, “h1” or “nc”
		Default	n/a

Table 103 **PCC Rule TLVs in the alc.sub_svc.pccrule Module (Continued)**

PCC Rule TLV	M|O	TLV Details	

pccrule.qos_action_ rate_limit_cir (Integer)	O (1)	Purpose	PCC rule action: rate-limit CIR
			• instantiate a dynamic policer for all flows in the PCC rule (if not already present)
			• set the CIR value
			Can be applied on ingress and egress.
			Results in IPv4 and/or IPv6 criteria entry in QoS policies
			CLI equivalent:
			`policer 1 # dynamic policer`
			`rate … cir <cir>`
			`exit`
			`entry 10 create`
			`match`
			`…`
			`exit`
			`action policer 1`
			`exit`

		Value	0 to 2000000000 kb/s
		Default	n/a

pccrule.qos_action_ rate_limit_pir (integer)	O (1)	Purpose	PCC rule action: rate-limit PIR
			• Instantiate a dynamic policer for all flows in the PCC rule (if not already present)
			• Set the PIR value
			Can be applied on ingress and egress.
			Results in IPv4 and/or IPv6 criteria entry in QoS policies.
			CLI equivalent:
			`policer 1 # dynamic policer`
			`rate <pir>`
			`exit`
			`entry 10 create`
			`match`
			`…`
			`exit`
			`action policer 1`
			`exit`

		Value	1 to 2000000000 kb/s
		Default	None

Table 103 **PCC Rule TLVs in the alc.sub_svc.pccrule Module (Continued)**

PCC Rule TLV	M|O	TLV Details	

pccrule.qos_action (integer)	O (1)	Purpose	PCC rule action: QoS forward
			Can be applied on ingress and egress
			Results in IPv4 and/or IPv6 criteria entry in QoS policies
			CLI equivalent:
			`entry 10 create`
			`  match`
			`  …`
			`  exit`
			`  action`
			`exit`

		Value	pccrule.qos_action_forward (1)
		Default	n/a

pccrule.filter_action_http_redirect (string)	O (1)	Purpose	PCC rule action: http-redirect
			Can be applied on ingress only
			Results in an IPv4 and/or IPv6 filter entry
			CLI equivalent:
			`entry 10 create`
			`  match next-header tcp`
			`  ...`
			`  exit`
			`  action`
			`  http-redirect <rdr-url-string>`
			`  exit`
			`exit`

		Value	http-redirect URL string (maximum 255 characters)
		Default	n/a

Table 103 **PCC Rule TLVs in the alc.sub_svc.pccrule Module (Continued)**

PCC Rule TLV	M|O	TLV Details	

pccrule.filter_action_redirect_to_nexthop_v4
(string)
	O (1)	Purpose	PCC rule action: redirect to a next-hop IPv4 address
			Can be applied on ingress only
			Results in an IPv4 filter entry
			CLI equivalent:
			`entry 10 create`
			`  match`
			`  ...`
			`  exit`
			`  action`
			`  forward next-hop <ip-address>`
			`  exit`
			`exit`

		Value	IPv4 address
		Default	n/a

pccrule.filter_action_redirect_to_nexthop_v6
(string)
	O (1)	Purpose	PCC rule action: redirect to a next-hop IPv6 address
			Can be applied on ingress only
			Results in an IPv6 filter entry. CLI equivalent:
			`entry 10 create`
			`  match`
			`  ...`
			`  exit`
			`  action`
			`  forward next-hop <ipv6-address>`
			`  exit`
			`exit`

		Value	IPv6 address
		Default	None

Table 103 **PCC Rule TLVs in the alc.sub_svc.pccrule Module (Continued)**

PCC Rule TLV	M|O	TLV Details	

pccrule.filter_action_ redirect_to_router_v4 (integer)	O (1)	Purpose	PCC rule action: redirect to a routing instance
			Can be applied on ingress only
			Results in an IPv4 filter entry
			CLI equivalent:
			`entry 10 create`
			`  match`
			`  ...`
			`  exit`
			`  action`
			`  forward router <router-instance>`
			`  exit`
			`exit`

		Value	service-id
		Default	n/a

pccrule.filter_action_ redirect_to_router_v6 (Integer)	O (1)	Purpose	PCC rule action: redirect to a routing instance
			Can be applied on ingress only
			Results in an IPv6 filter entry
			CLI equivalent:
			`entry 10 create`
			`  match`
			`  ...`
			`  exit`
			`  action`
			`  forward router <router-instance>`
			`  exit`
			`exit`

		Value	*service-id*
		Default	n/a

Table 103 **PCC Rule TLVs in the alc.sub_svc.pccrule Module (Continued)**

PCC Rule TLV	M|O	TLV Details	

pccrule.filter_action (Integer)	O (1)	Purpose	PCC rule action: Filter forward or drop
			Can be applied on ingress and egress
			Results in an IPv4 and/or IPv6 filter entry
			CLI equivalent:
			`entry 10 create`
			`  match`
			`  ...`
			`  exit`
			`  action`
			`  forward | drop`
			`  exit`
			`exit`

		Value	pccrule.filter_action_forward (1) pccrule.filter_action_drop (2)
		Default	n/a

Notes:

- (1) At least one PCC rule action must be specified.
- M=Mandatory, O=Optional

Table 104 lists the alc.sub_svc.flow function.

Table 104 **alc.sub_svc.flow Functions**

alc.sub_svc.flow Functions	Description
sub_svc.flow.add_to_flow (flow, flow TLV, value)	Appends a flow TLV such as dscp, protocol, src-ip, dst-ip, src-port, or dst-port to the flow list. The flow list defines matching criteria for an IP flow and can be added to a PCC rule with the sub_svc.pccrule.add_to_pccrule() function. Parameters: flow (type = list): list containing the match criteria (DSP, 5-tuple) that describes an IP flow. Flow TLVs are appended to this list with the sub_svc.flow.add_to_flow() function. The flow is added to a PCC rule with the sub_svc.pccrule.add_to_pccrule() function. flow TLV (type int): Flow TLV that will be appended to the flow list. value (type as defined for the flow TLV): the value of the flow TLV that will be appended to the flow list

Table 105 lists the PCC Rule TLVs alc.sub_svc.flow Module.

Table 105 **PCC Rule TLVs alc.sub_svc.flow Module**

Flow TLV	M|O	TLV Details	
flow.dscp (string)	O	Purpose	Specifies a DSCP flow match criterion
		Value	Fixed DSCP name strings as in the output of **show qos dscp-table; f**or example, “be” or “ef”. The DSCP name must be specified in lower case.
		Default	n/a
flow.protocol (integer)	O	Purpose	Specifies a protocol number match criterion
		Value	0 to 255
		Default	n/a
flow.dst-ip (string)	O	Purpose	Specify a destination IPv4 or IPv6 match criterion
		Value	*ipv4-address* | *ipv6-address* | any where *ipv4-address*: d.d.d.d[/m] d [0 to 255] m [0 to 32] *ipv6-address*: x:x:x:x:x:x:x:x[/preflen] x: [0 to FFFF] preflen: 0 to 128
		Default	any
flow.dst-port (string)	O	Purpose	Specify a destination port match criterion
		Value	port or port range: *port*[-*port*] where *port*: 0 to 65535
		Default	n/a

Table 105 **PCC Rule TLVs alc.sub_svc.flow Module (Continued)**

Flow TLV	M|O	TLV Details	
flow.src-ip (string)	O	Purpose	Specify a source IP or IPv6 match criterion
		Value	*ipv4-address* | *ipv6-address* | any where *ipv4-address*: d.d.d.d[/m] d [0 to 255] m [0 to 2] *ipv6-address*: x:x:x:x:x:x:x:x[/preflen] x: [0 to FFFF] preflen: 0 to 128
		Default	any
flow.src-port (string)	O	Purpose	Specifies a source port match criterion
		Value	port or port range: *port*[-*port*] where *port*: 0 to 65535
		Default	n/a

9.18.9 Operational Commands

9.18.9.1 Show Commands

To display the active subscriber services in the system, use the **show service sub-services** CLI command. The **sub-service-name** filter is a longest match.

```
# show service sub-services [id <service-id>] [sap <sap-id>] [ip <ip-prefix/prefix-
length>] [mac <ieee-address>] [pppoe-session-id <pppoe-session-id>] [sub-service-
name <sub-service-name>] [sub-service-type <sub-service-
type>] [summary|associations]
```

Sample output:

```
# show service sub-services
===============================================================================
Subscriber service table
===============================================================================
Username        : cpe2@domain1.com
Subscriber      : cpe2@domain1.com                    Type     : PPP
SAP             : 1/1/4:1201.5                        Service  : 1000
```

```
MAC Address     : 00:51:00:00:01:01                    PPPoE-SID: 114
IP              : 10.1.1.3
Interface-ID    : 02:51:00:FF:FE:00:01:01
-------------------------------------------------------------------------------
  Service instance : rate-limit;5120;30720
  up time          : 0d 00:05:23
  type             : rate-limit
  acct sess id     : 144DFF00000A2556B757DA
  multi sess id    : 144DFF00000A1156B71E44
  acct type        : volume-time
  acct ivl         : 1d 00:00:00
  input octets     : 0
  input packets    : 0
  output octets    : 0
  output packets   : 0
  actions          :
  QoS-override
    ingress arbiter "root" rate 5120
  QoS-override
    egress aggregate rate limit 30720
-------------------------------------------------------------------------------
number of subscriber services found: 1
===============================================================================
```

To display the active PCC rules in the system, use the **show service active-subscribers pcc-rule** CLI command. A PCC rule can be inactive when, for example, a PCC rule with filter actions on IPv6 flows is activated on an IPv4single-stack PPPoE or IPoE session.

```
# show service active-subscribers pcc-rule [subscriber <sub-ident-string>] [detail]
```

Sample output:

```
# show service active-subscribers pcc-rule subscriber "ipoe-bridged-001" detail
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber ipoe-bridged-001 (sub-profile-1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[1/1/4:1201.19] - sla:sla-profile-1
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
2001:db8:1:101::1/128
                00:51:00:00:00:05    IPoE           DHCP6        1000       Y
-------------------------------------------------------------------------------
Ingr Filter Override     : N/A
Egr  Filter Override     : N/A
Ingr Qos Policy Override : 10:P10
Egr  Qos Policy Override : 10:P9
===============================================================================
Precedence Rule Id   Rule Name                Subscriber Service Name
-------------------------------------------------------------------------------
```

```
10         19        rule-1                   subsvc_pcc-coa-1
10         20        rule-2                   subsvc_pcc-coa-1
===============================================================================
===============================================================================
PCC Rules
===============================================================================
PCC rule name         : rule-1
PCC rule id           : 19
Monitoring key        : -
Flow status           : Enabled
Nbr of Flows          : 2 (ingress)
HTTP-Redirect         : -
Next-Hop Redir. IPv4  : -
Next-Hop Redir. IPv6  : -
QoS Ingr. CIR/PIR     : - / 1000 kbps
QoS Egr. CIR/PIR      : - / -
FC change             : -
Account               : Enabled
-------------------------------------------------------------------------------
Flows
-------------------------------------------------------------------------------
Src. IP  : any                                 Src. Port: -
Dst. IP  : 172.16.1.1/32                       Dst. Port: -
Protocol : -                                   DSCP     : -
-------------------------------------------------------------------------------
Src. IP  : any                                 Src. Port: -
Dst. IP  : 2001:db8:aaa:1::1/128               Dst. Port: -
Protocol : -                                   DSCP     : -
-------------------------------------------------------------------------------
===============================================================================
PCC rule name         : rule-2
PCC rule id           : 20
Monitoring key        : -
Flow status           : Enabled
Nbr of Flows          : 2 (egress)
HTTP-Redirect         : -
Next-Hop Redir. IPv4  : -
Next-Hop Redir. IPv6  : -
QoS Ingr. CIR/PIR     : - / -
QoS Egr. CIR/PIR      : - / 5000 kbps
FC change             : -
Account               : Enabled
-------------------------------------------------------------------------------
Flows
-------------------------------------------------------------------------------
Src. IP  : 172.16.1.1/32                       Src. Port: -
Dst. IP  : any                                 Dst. Port: -
Protocol : -                                   DSCP     : -
-------------------------------------------------------------------------------
Src. IP  : 2001:db8:aaa:1::1/128               Src. Port: -
Dst. IP  : any                                 Dst. Port: -
Protocol : -                                   DSCP     : -
-------------------------------------------------------------------------------
===============================================================================
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
```

Use the following alternative command to check the PCC rules in the system:

```
# show subscriber-mgmt pcc-rule
  - pcc-rule
  - pcc-rule monitoring-key <key> detail
  - pcc-rule rule-id <id> detail
  - pcc-rule rule-name <rule-name>
  - pcc-rule rule-name <rule-name> detail
  - pcc-rule summary
  - pcc-rule monitoring-key <key>
 <rule-name>          : [100 chars max]
 <key>                : [80 chars max]
 <id>                 : [1..1023]
 <detail>             : keyword
```

The statistics of dynamic policers can be displayed with:

```
# show service active-subscribers subscriber "ipoe-bridged-001" detail
--- snip ---
-------------------------------------------------------------------------------
SLA Profile Instance per Policer statistics
-------------------------------------------------------------------------------
                        Packets                 Octets
--- snip ---
Ingress Policer 10 (Stats mode: minimal)
used by pcc-rule rule-1
Off. All              : 0                       0
Dro. All              : 0                       0
For. All              : 0                       0
--- snip ---
```

The details of the cloned QoS and filter policies as a result of PCC rule activation can be displayed with the following show commands:

```
# show qos sap-ingress 10:P1
# show qos sap-egress 10:P1
# show filter ip 10:P1
# show filter ipv6 "10:P1"
```

9.18.9.2 Debug Commands

There are no specific RADIUS subscriber services debug commands. The debugging is part of the RADIUS and Python debug output; for example:

```
debug
    router "Base"
        radius
            packet-type authentication accounting coa
            detail-level high
        exit
    exit
    python
        python-script "subsvc-1"
            script-all-info
        exit
```

```
    exit
exit
```

9.18.9.3 Resource Monitoring

For information on resource monitoring, see PCC Rules and Capacity Planning and PCC Rule Scaling Example.

The following CLI command provides an overview of the resource usage per line card, such as the number of ACL and ACL QoS entries, Filters, QoS policies, dynamic policers, and QoS overrides:

```
# tools dump resource-usage card [<slot-number>] all
```

These resource counters are available in SNMP and can be used in RMON to trigger threshold crossing alarms; for example:

```
configure system
    thresholds
        rmon
            alarm 1 variable-
oid tFPResIngIPv6AclEntryAlloc.1.1.1 interval 10 rising-event 1 rising-
threshold 25000 falling-event 2 falling-threshold 24000
            event 1 description "Ingress IPv6 ACL Entries too high"
            event 2 description "Ingress IPv6 ACL Entries - below limit"
        exit
```

The summary output of the **show subscriber-mgmt pcc-rule** command lists the number of active PCC rules and the number of active combinations:

```
# show subscriber-mgmt pcc-rule summary
===============================================================================
PCC Rules Summary
===============================================================================
Total Nbr PCC Rules     : 2 / 1024
Nbr Active PCC Rules    : 2 / 1024
Nbr Active Combinations
  IPv4 Filter           : 0 / 4095
  IPv6 Filter           : 0 / 4095
  Egress Qos            : 1 / 4095
  Ingress Qos           : 1 / 4095
===============================================================================
```

9.19 Residential Gateway Replacement

Residential gateway (RG) replacements are performed for a variety of reasons such as upgrading hardware, replacing broken equipment, or relocating to a new home. However, the BNG's anti-spoof filter and host-limit features can sometime prevent immediate RG replacement. In some cases, a subscriber must wait for an old DHCP lease to expire before a new RG is allowed to connect to the BNG. For example, some service providers assign an IP address and/or prefix based on physical line, SAP-id, circuit-id, or interface-id. Therefore, a home is always assigned the same IP address and/or prefix. On the BNG, an anti-spoof mechanism prevents different MAC addresses from using the same IP address. As a result, the new RG fails the anti-spoof filter and is denied an IP address and/or prefix. The subscriber in this case must wait for the DHCP lease of the RG to time out in order for the anti-spoof filter to remove its entry.

Two features, **lease-override** and **shcv-policy**, may help improve the RG replacement process. These features focus on minimizing service interruption and enhancing the end subscriber experience. RGs, in general, have no mechanisms to inform the BNG or the DHCP server that they have been disconnected from the network. Even if the BNG has periodic SHCV enabled, the detection might take some time. Often, when a subscriber plugs in a new RG, the BNG still has the old RG registered as a host. This has two consequences. First, if the new RG is assigned the same IP address as the old RG, then an IP-conflict occurs and fails the anti-spoof filter. Second, if the SAP has a host limit or a session limit provisioned, then exceeding the limit will prevent the new RG from receiving an IP address and/or prefix.

Starting in Release 13.0R4, if an IP conflict occurs on the same SAP, then by default the new RG (MAC) will immediately override the DHCP lease of the old device. This is known as **lease-override**. This is applicable to DHCP relay and proxy for both IPv4 and IPv6 hosts. Prior to Release 13.0.R4, lease-override only occurred for DHCPv4 relay. The **lease-override** is performed only when an IP-conflict occurs within the same SAP.

The other option is to use trigger SHCV to check the connectivity status of the old RG before removing it and its lease. This is known as the ip-conflict-triggered SHCV under the SHCV-policy. The SHCV is sent only when the BNG detects an IP address conflict on DHCP discovery. If the host does not respond within the configured timeout, both the host and lease are removed from the BNG. The new RG is required to perform a subsequent DHCP discovery or request in order to install a host. SHCV can help prevent malicious RGs from spoofing another RG IP address. Trigger SHCV for IP-conflict is available for DHCPv4/v6 relay and proxy, as well as ARP hosts. The following table specifies when the SHCV is sent for IP-conflict.

Table 106 **IP-Conflict SHCV Triggering Points**

Configuration on group interface	Triggered on
DHCPv4 proxy	DHCP Discovery
DHCPv4 relay	DHCP Request
DHCPv6 proxy	DHCP Advertisement
DHCPv6 relay	DHCP Advertisement
ARP-host	Host's initial ARP

It is also possible that new RGs are denied of service because of a host limit set against the subscriber. There are various host limit features on the BNG: sla-profile host-limits, ipoe session-limit, and ipoe sap-session limit. For example, setting a host limit of "1" can ensure that each home will only take one IP address. As mentioned before, RGs do not inform the BNG of a disconnect. If SHCV is enabled, it might take some time before the BNG is informed of the disconnect. Therefore, when a new RG connects to the BNG, the BNG will perform a host limit check (if configured) against the subscriber. If the old host still has an entry on the BNG and there is a host-limit of "1", the new RG is denied an IP address and/or prefix assignment because it has exceeded the host limit. A trigger SHCV, "host-limit-exceeded" inside the shcv-policy can be configured against a group-interface. This SHCV is triggered when an over limit is detected. If the existing host registered on the BNG does not respond within the configured timeout both the host and its lease are removed from the BNG. The SHCV can only remove hosts from the BNG, the new RG is still required to perform a subsequent DHCP discoveries or requests to obtain an IP address.

By providing lease-override and various SHCV triggers in the SHCV policy, service providers have a variety of options to allow subscribers to perform quick and seamless RG replacements.

It is possible to use host-connectivity check without the shcv-policy. The main function of host-connectivity check would be for periodic check. The trigger functions are performed through the SHCV-policy.

9.20 ESM Troubleshooting Show Command

Network operators are sometimes unable to turn on debug to troubleshoot customers issues on a live production network. Turning on debug might affect the BNG performance, and some support technicians may not have access to debug and configuration commands.

The **show subscriber-mgmt errors** command is a show command that captures detailed ESM errors. This command helps diagnose problems immediately without the need to turn on the debug function. Only DHCPv4 and PPPv4 supports this command; some support details are provided below. This command will not compromise the BNG performance and does not require debug or configuration commands.

- DHCPv4 support:
  - dropped protocol messages (host setup/renew)
  - protocol timeouts
- PPPv4 support (including PPPoE, PPPoA, and PPPoEoA):
  - dropped protocol messages
  - traps

IPv6 host setup errors may be captured in error logs.

All subscriber problems are first stored in a main circular buffer. The main circular buffer is then fed into smaller circular buffers, organized by MAC addresses. When the buffer is full, the first circular buffer will purge the oldest message to make room for the newest message. The smaller circular buffers (one per MAC address) store a limited number of messages per MAC. Again, the smaller buffer will delete the oldest to make room for the newest. The circular buffer, per MAC, prevents one device from holding all the error messages in the buffer. The main circular buffer can hold 5,000 errors in total, while the smaller buffer can hold 10 log entries per MAC. The circular buffer supports CPM3 and higher.

The **show** command allows sorting by MAC, subscriber SAP, SDP, and unknown-origin (unknown SAP or SDP). The **show** command allows the input of a specific MAC, SAP, or SDP to directly search for particular subscriber issues or problems.

The circular buffer will only log drop reasons for DHCPv4 and PPPoEv4. Non-error reasons are never logged; for example, a drop due to being in SRRP standby is not logged. The circular buffer has a timestamp associated with each error and the errors are listed beginning with the most recent. Error logs are lost on a HA switchover, and persistency is not supported. There is no throttling mechanism for the same errors; it is possible to fill the circular buffer with the same error message from different MAC addresses.

9.21 Subscriber Accumulated Statistics

The accumulated statistics policy defines which statistics for queues and policers of a particular subscriber should be collected and stored in memory. At the end of the subscriber session, the queue and policer statistics are added to the statistics already stored in memory from previous sessions. This enables operators to view queue and policer statistics even if the subscriber is offline.

The accumulated statistics policy supports up to four ingress and four egress entries. For queue statistics, v4-v6 mode is not supported, where v4 and v6 statistics are always aggregated. For policer statistics, only min-mode is supported.

When a single subscriber has a list of bridge hosts, all hosts are forced to use the same statistics policy. If hosts use a different SLA profile and the operator wants to collect the statistics for all hosts, the statistics policy must encompass all queues and policers for various SLA profiles. If there are multiple SLA profile instances for the same subscriber, the statistics are summed up for each instance on a per policer or queue basis. These statistics are not exchanged between MCS peers. Therefore, for dual-homed hosts, the statistics need to be gathered from all the nodes and then summed up. If a queue or a policer is missing from the accumulated statistics policy in the current subscriber session and offline statistics exist for that entity from previous sessions, these offline stats will be lost when the current subscriber session ends.

The **show subscriber-mgmt accu-stats subscriber** command displays the cumulative statistics for a subscriber. If the subscriber is online, the cumulative statistics consist of the sum of both the subscriber statistics while online and the offline statistics. If the subscriber is offline, the last subscriber session statistics are added to the offline statistics.

When the accumulated statistics policy changes through a subscriber profile override, either through the **tools** command or through RADIUS, the stored statistics can be affected. If the new SLA profile uses the same queues and policers already stored in memory, these statistics continue to accumulate. Only queues and policers that differ are added and start from zero.

If resources for capturing offline statistics are full, a trap is generated in log 99 to warn the operator. The command **show subscriber-mgmt status system** shows the number of subscribers using these accumulated statistics and a flag in this command shows whether the usage is at its peak value.

In the **show subscriber-mgmt accu-stats-subscribers** command, active subscribers with an accumulated statistics policy configured have both the sub-profile-name and accu-stats-policy fields populated. Active subscribers that are no longer associated with an accumulated statistics policy have accu-stats-policy populated with “Unknown”. Inactive subscribers have “Unknown” for both the sub-profile-name and the accu-stats-policy fields. This command is useful when combined with the CLI **match** command, to search for subscribers with certain properties. For example, **match “Unknown”** displays the list of subscribers that are no longer associated to any accumulative statistics policy.

It is possible for an active subscriber to have offline statistics without an accumulated statistics policy if the accu-stats-policy was removed from the sub-profile.

When the active subscriber has an accumulated statistics policy, the subscriber’s offline statistics can be deleted using the following command.

clear subscriber-mgmt accu-stats subscriber *subscriber-id*

The **show subscriber-mgmt accu-stats subscriber** *subscriber-id* command then displays only the subscriber’s current statistics as no statistics remain in the offline storage.

If an active subscriber does not have an accumulated statistics policy, the subscriber’s offline statistics can be deleted in one of the following ways.

To remove the offline statistics for all active subscribers that are no longer associated with an accumulated statistics policy, the following command can be used.

clear subscriber-mgmt accu-stats active-subs no-accu-stats-policy

To remove the offline statistics for a group of active subscribers that is no longer associated to an accumulated statistics policy and that has a defined subscriber profile (for example, if the accumulated statistics policy has been removed from the subscriber profile), the following command can be used.

clear subscriber-mgmt accu-stats active-subs sub-profile *profile-name*

It is also possible to remove the offline statistics for an inactive (offline) subscriber. To remove offline statistics for all inactive subscribers, use the following command.

clear subscriber-mgmt accu-stats inactive-subs

To remove offline statistics for a specific inactive subscriber, use the following command.

clear subscriber-mgmt accu-stats subscriber *subscriber-id*

The subscriber accumulated statistics feature is supported only on systems with CPM-5 and up, SR-e, and SR-A.

9.22 Hybrid Access

In a hybrid access deployment, a home's residential gateway (RG) is connected to the network via both a wired and a wireless link. Traffic can be split over these links via various mechanisms such as per-flow hashing, flow policies, or MP-TCP. Both access connections must be terminated in a common endpoint called the Hybrid Access Gateway (HAG). This gateway provides Internet connectivity to the home.

9.22.1 BNG-based HAG

SR OS TPSDA supports hybrid access deployments where the BNG acts as a Serving Gateway (SGW), PDN Gateway (PGW), and HAG. In this model, both the fixed access and wireless access links share the same Layer 3 (L3) IP/IPv6 address. The RG/HAG determines which Layer 2 (L2) connection should be used. To support this model, SR OS supports GTP termination and ESM connection bonding. The BNG is the L3 gateway in this model and attracts all IP traffic. Multicast traffic is supported.

Figure 167 shows a sample Hybrid Access deployment with a BNG-based HAG.

Figure 167 **SGW/PGW/BNG Integrated Hybrid Access Gateway**

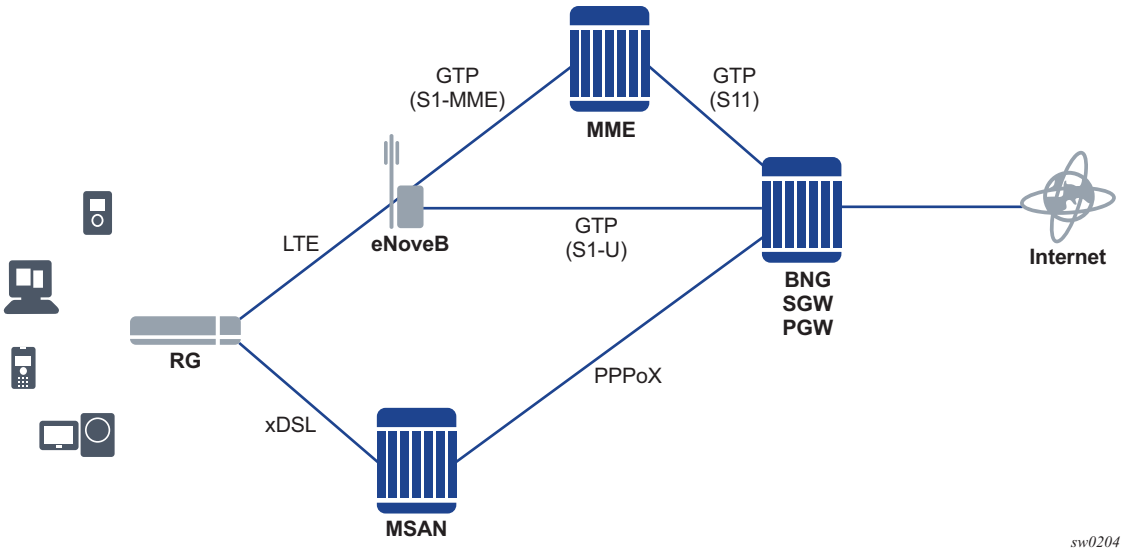

9.22.2 PGW-based HAG

SR OS TPSDA supports hybrid access deployments where a PGW acts as a HAG. In this model, both the fixed access and wireless access links share the same L3 IP/IPv6 address. The RG/HAG determines which L2 connection should be used. To support this model, SR OS IPoE session and PPPoE session functionality is extended to connect to a GTP uplink. When this uplink is active, all unicast traffic is forwarded via a GTP tunnel between the BNG and PGW. This way, the PGW acts as the L3 gateway and the BNG does not attract subscriber traffic via regular routing. Multicast traffic is not supported and should be handled out-of-band.

Figure 168 shows a sample Hybrid Access deployment with a PGW-based HAG.

Figure 168 **Sample Hybrid Access Deployment**

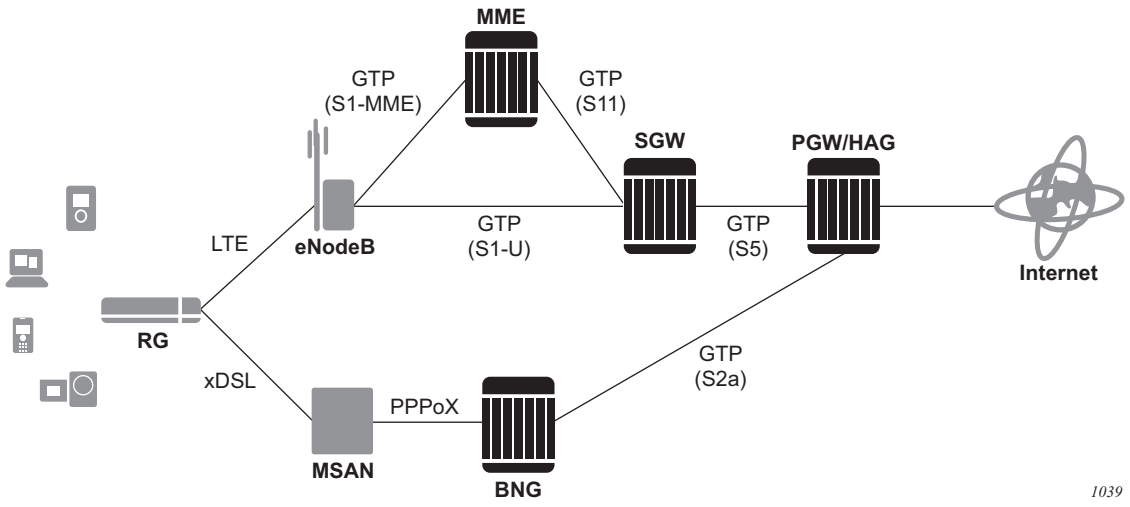

For more details on GTP uplinks, see the GTP section.

9.23 Connection Bonding

ESM connection bonding allows two Layer 2 access connections, for example, GTP and PPPoE, to combine to share a single Layer 3 IP connection. This allows the seamless use of two connections for additional bandwidth or resiliency without the need to manage multiple IP addresses.

SR OS spreads downstream traffic, either based on fixed weights, a filter decision, a pure active/standby, or dynamic load-balancing that attempts to saturate one link before using another link. Upstream traffic can enter through either connection, but it is recommended to keep flows identified by 5-tuple on the same link to avoid reordering.

Connection bonding requires an FPE type of **sub-mgmt-extensions**.

9.23.1 Setup

During normal authentication, access connections can indicate they are part of a common bonding context by specifying a bonding identifier. When the first connect is set up, an additional authentication phase is started for the bonding context itself. Figure 169 shows RADIUS-based authentication for bonding of an IPoE and GTP access connection. For simplicity, all access nodes, such as the residential gateway, MSAN, eNodeB, and MME have been identified as a single entity.

Figure 169 **Bonding authentication**

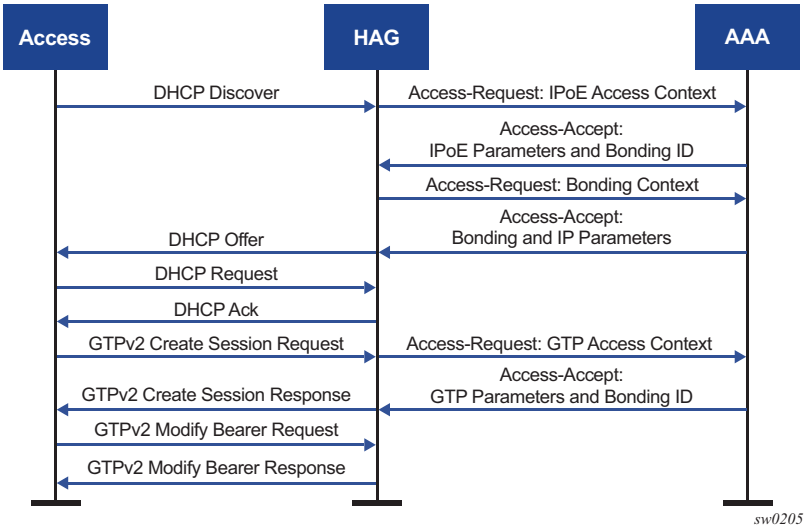

All address assignments and Layer 3 parameters are shared between the access connections and are therefore handled in the bonding context. The system supports either Local Address Assignment or AAA-provisioned IP addresses. Access connections cannot use any DHCPv6 relay or client mechanisms.

After the setup is complete, ESM subscriber resources are created for each context as follows.

- a unique subscriber with a single internal IPoE session to represent the bonding context

  This is the main context for functionality, including QoS, filters, and accounting. This subscriber cannot be reused for any other hosts, regardless if they are bonding or not.
- a subscriber per access connection with one or more hosts as needed, depending on the access type

  This subscriber must be distinct from the bonding subscriber. Other non-bonded hosts or sessions may be present under the same subscriber. A subset of ESM features (for example, QoS, filters, and accounting) are also available in this context; however, it is recommended that the majority of the feature be configured in the bonding context.

All access and bonding ESM contexts need to be present in distinct VRFs. The bonding context must be created in a special group interface of type **bonding**. This group interface is not linked to any SAPs, however, an FPE construct ensures the link between the access and bonding context.

9.23.2 Downstream Load-Balancing

By default, downstream packets are hashed over two connections on a per-flow basis, allowing packets of the same flow to follow the same path and avoiding reordering issues. Flows are identified by the 5-tuple <src-ip, dst-ip, protocol, src-port, dst-port>. The hash weights of each connection are configurable. It is possible to override the hash-based connection selection with an IP-filter based selection.

The initial hash weights can also be dynamically adapted based on the load on the primary connection. The IOM periodically measures how much traffic is sent over the primary connection, comparing it to a predefined saturation rate. When the connection is saturated, the hash weights automatically change to send more traffic over the alternate link. Similarly, if the total rate of traffic decreases, the hashing weights change to send more traffic over the primary connection.

For dynamic load-balancing, the following must be defined:

- the primary connection, identified as connection 1 in the bonding context
- the reference rate in the primary connection, which defines the saturation point

  By default, the aggregate rate is used, but this can be changed during AAA authentication. The saturation point is further refined using thresholds that define when the rate change should be executed.
- a weight change value, specifying how much the weight percentages should change each time the saturation point is reached
- policers for all downstream traffic in the bonding context

  A single policer that directly feeds into multiple local queues can be used.

When only a single connection is active, all traffic is sent to this connection, regardless of hashing weights or filters.

9.23.3 QoS

Regular ESM QoS is supported in both the access and bonding contexts; however, there is no direct feedback mechanism between the two contexts. Therefore, if an access connection drops a packet, it is not reflected in bonding statistics, nor does it cause backpressure on the bonding QoS algorithm.

When traffic passes over the FPE from the bonding context to the access context or from the access context to the bonding context, the system keeps the traffic classification and the in- and out-profile markings. Although this occurs automatically, bonding subscriber policies for ingress and egress should consider the following recommendations.

- Enable **de-mark** for access egress and map each FC to the dot1p as defined in Table 107, thus ensuring that the same classification is used in the access connection context.
- Perform classification for access ingress based on dot1p as defined in Table 107 and enable **in-profile** and **de-1-out-profile** for each FC, thus ensuring the same classification is used as for the access connection context. A different classification scheme can be used if required, for example, based on DSCP or IP criteria.

Table 107 **FC to dot1p Mapping**

FC	dot1p
be	0
l2	1

Table 107 **FC to dot1p Mapping (Continued)**

FC	dot1p
af	2
l1	3
h2	4
ef	5
h1	6
nc	7

To support load-balancing in the bonding context, the configured **stat-mode** of any egress policer in the bonding context is ignored. Instead, an internal stat-mode is used, which uses two counters (one per access connection), which is reflected in the in- and out-of-profile statistics.

9.23.4 Multicast

Multicast replication is supported in context of the access connections. By default, multicast streams are replicated in the connection where the corresponding IGMP/MLD join is received; however, this can be overridden to always force a specific connection.

When one connection fails, multicast replication automatically sets up in the alternate connection and does not require a new IGMP/MLD packet to arrive.

MCAC bandwidths must be configured equally on both access connections; otherwise, there may be unexpected drops of (S,G) pairs.

If a multicast stream is forwarded over the primary connection and **egress-rate-modify** is in use, any potential change of the reference rate is taken into consideration by the load-balancing mechanism for unicast traffic as described in Downstream Load-Balancing. When using per-host replication for a bonded host, similar adaptations are made based on channel definitions in applicable MCAC policies.

9.24 Configuring Enhanced Subscriber Management with CLI

This section provides information to configure subscriber management features using the command line interface. It is assumed that the reader is familiar with VPLS and IES services.

9.24.1 Configuring RADIUS Authentication of DHCP Sessions

When RADIUS authentication for subscriber sessions is enabled, DHCP messages from subscribers are temporarily held by the BSA, while the user's credentials are checked on a RADIUS server.

Configuring RADIUS authentication for subscriber sessions is done in two steps:

- First define an authentication-policy in the **config**>**subscriber-mgmt**>**authentication-policy** context.
- Then apply the policy to one or more SAPs in the **config**>**service**>**vpls**>**sap**>**authentication-policy** *auth-plcy-name* context (for a VPLS service).

  Or apply the policy to one or more interfaces **config**>**service**>**ies**>**if**>**authentication-policy** *auth-plcy-name* context (for an IES service):

The following example displays a partial BSA configuration with RADIUS authentication:

```
A:ALA-1>config>service# info
----------------------------------------------
subscriber-management
    authentication-policy BSA_RADIUS create
        description "RADIUS policy for DHCP users Authentication"
        password "mysecretpassword"
        radius-authentication-server
            server 1 address 10.100.1.1 secret "radiuskey"
            retry 3
            timeout 10
        exit
        re-authentication
        user-name-format circuit-id
    exit
exit
...
vpls 800 customer 6001
```

```
        description "VPLS with RADIUS authentication"
        sap 2/1/4:100 split-horizon-group DSL-group create
            authentication-policy BSA_RADIUS
        exit
        sap 3/1/4:200 split-horizon-group DSL-group create
            authentication-policy BSA_RADIUS
        exit
        no shutdown
    exit
    ...
    ----------------------------------------------
    A:ALA-1>config>service#
```

9.24.2 TCP MSS adjustment for ESM Hosts

TCP MSS adjustment is supported to prevent fragmentation of TCP packets from/to ESM hosts. See the TCP MSS Adjustment for ESM Hosts section of the Multiservice Integrated Service Adapter Guide.

9.24.3 Configuring Enhanced Subscriber Management

9.24.3.1 Basic Configurations

Configuring and applying the Enhanced Subscriber Management profiles and policies are optional. There are no default Profiles or policies.

The basic Enhanced Subscriber Management profiles and policies must conform to the following:

- Unique profile or policy names (IDs)
- Profiles and/or policies must be associated with a VPLS or IES service to facilitate Enhanced Subscriber Management.
- QoS and IP filter entries configured in Enhanced Subscriber Management profiles and policies override the defaults and/or modified parameters or the default policies.
- The Enhanced Subscriber Management profiles and policies must be configured within the context of VPLS or IES.

9.24.3.2 Subscriber Interface Configuration

The following output displays a basic subscriber interface configuration.

```
*A:ALA-48>config>service>ies>sub-if# info
----------------------------------------------
                description "Routed CO - Antwerp 2018"
                address 192.168.2.254/24
                address 192.168.3.254/24
                address 192.168.4.254/24
                address 192.168.5.254/24
                address 192.168.6.254/24
                group-interface "DSLAM_01" create
                    description "Routed CO - vlan / subscriber"
                    sap 1/1/2:1001 create
                        static-host ip 192.168.2.2 create
                        exit
                    sap 1/1/2:1002 create
                        static-host ip 192.168.2.2 create
                        exit
                    sap 1/1/2:1004 create
                        static-host ip 192.168.2.4 create
                        exit
                    sap 1/1/2:1100 create
                        static-host ip 192.168.2.100 create
                        exit
                    exit
                exit
----------------------------------------------
*A:ALA-48>config>service>ies>sub-if#
```

9.24.3.3 Configuring Enhanced Subscriber Management Entities

9.24.3.3.1 Configuring a Subscriber Identification Policy

The following displays an example of a subscriber identification policy configuration:

```
A:ALA-48>config>subscr-mgmt# info
----------------------------------------------
...
        sub-ident-policy "Globocom" create
            description "Subscriber Identification Policy Id Globocom"
            sub-profile-map
                entry key "1/1/2" sub-profile "ADSL Business"
            exit
            sla-profile-map
                entry key "1/1/2" sla-profile "BE-Video"
            exit
            primary
                script-url "primaryscript.py"
```

```
                no shutdown
            exit
            secondary
                script-url "secundaryscript.py"
            exit
            tertiary
                script-url "tertiaryscript.py"
                no shutdown
            exit
        exit
...
-------------------------------------------------
A:ALA-48>config>subscr-mgmt#
```

9.24.3.3.2 Configuring a Subscriber Profile

Enhanced Subscriber Management subscriber profile configurations specify existing QoS scheduler profiles. In the following example, “BE-Video-max100M” is specified in the sub-profile “ADSL Business” for the ingress-scheduler-policy. “Upload” is specified in the sub-profile egress-scheduler-policy.

```
#--------------------------------------------------
echo "QoS Policy Configuration"
#--------------------------------------------------
    qos
        scheduler-policy "BE-Video-max100M" create
            description "Scheduler Policy Id BE-Video-max100M"
            tier 1
                scheduler "tier1" create
                    description "Scheduler Policy Id BE-Video-max100M Tier 1 tier1"
                exit
            exit
        exit
        scheduler-policy "Upload" create
            description "Scheduler Policy Id Upload"
            tier 3
                scheduler "tier3" create
                    description "Scheduler Policy Id Upload Tier 3 tier3"
                exit
            exit
        exit
        sap-ingress 2 create
            description "Description for Sap-Ingress Policy id # 2"
            queue 1 create
            parent "tier1"
            exit
            queue 11 multipoint create
            parent "tier1"
            exit
        exit
        sap-egress 3 create
            description "Description for Sap-Egress Policy id # 3"
            queue 1 create
            parent "tier3"
            exit
```

```
        exit
    exit
#-----------------------
```

The following displays an example of a subscriber identification policy configuration:

```
A:ALA-48>config>subscr-mgmt# info
----------------------------------------------
...
        sub-profile "ADSL Business" create
            description "Subscriber Profile Id ADSL Business"
            ingress-scheduler-policy "BE-Video-max100M"
                scheduler "tier1" rate 99
            exit
            egress-scheduler-policy "Upload"
                scheduler "tier3" rate 1 cir 1
            exit
            sla-profile-map
                entry key "1/1/3" sla-profile "BE-Video"
            exit
        exit
----------------------------------------------
A:ALA-48>config>subscr-mgmt#
```

9.24.3.3.3 Configuring an SLA Profile

The following displays an example of a SLA Profile configuration:

```
A:ALA-48>config>subscr-mgmt# info
------------------------------------------------
    subscriber-mgmt
        sla-profile "BE-Video" create
            description "SLA Profile Id BE-Video"
            ingress
                qos 2
                    queue 1
                    exit
                exit
            exit
            egress
                qos 3
                    queue 1
                    exit
                exit
            exit
        exit
----------------------------------------------
A:ALA-48>config>subscr-mgmt#
```

9.24.3.3.4 Configuring Explicit Mapping Entries

The following displays an example of an explicit subscriber mapping:

```
A:ALA-7>config>subscr-mgmt# info
----------------------------------------------
A:ALA-48>config>subscr-mgmt# info
----------------------------------------------
...
        explicit-subscriber-map
            entry key "1/1/1:1111" sub-profile "ADSL GO" alias "Sub-Ident-1/1/1:
1111" sla-profile "BE-Video"
        exit
...
----------------------------------------------
A:A:ALA-48>config>subscr-mgmt#
```

9.24.3.4 Routed CO with Basic Subscriber Management Features

The following displays the output of an IES service configured with and without enhanced subscriber management and only applies to the 7750 SR.

```
A:term17>config>service>ies# inf
----------------------------------------------
            subscriber-interface "s2" create
                address 10.20.1.1/16
                dhcp
                    gi-address 10.20.1.1
                exit
                group-interface "g3" create
                    description "With Enhanced Subscriber Mgmt"
                    arp-populate
                    dhcp
                        server 10.1.1.1
                        trusted
                        lease-populate 8000
                        no shutdown
                    exit
                    sap lag-1:11 create
                        sub-sla-mgmt
                            def-sub-profile "subProf"
                            def-sla-profile "slaProf"
                            sub-ident-policy "foo"
                            multi-sub-sap
                            no shutdown
                        exit
                        host ip 10.20.1.10 mac 00:00:aa:aa:aa:dd subscriber "One" su
b-
profile "subProf" sla-profile "slaProf"
                        exit
                    exit
            exit
            subscriber-interface "s3" create
                address 10.39.1.1/16
```

```
                    dhcp
                        gi-address 10.39.1.1
                    exit
                    group-interface "g5" create
                        description "Without Enhanced Subscriber Mgmt"
                        arp-populate
                        dhcp
                            server 10.1.1.1
                            trusted
                            lease-populate 8000
                            no shutdown
                        exit
                        sap 4/1/1:24.4094 create
                        exit
                    exit
                exit
                no shutdown
----------------------------------------------
A:term17>config>service>ies#
```

9.24.3.5 Applying the Profiles and Policies

Note: Subscriber interfaces operate only with basic (or enhanced) subscriber management. At the very least, a host, either statically configured or dynamically learned by DHCP must be present in order for the interface to be useful. This note applies to the 7750 SR only.

Apply the Enhanced Subscriber Management profiles and policies to the following entities:

- SLA Profile
- Subscriber Identification Policy
- Subscriber Profile

9.24.3.5.1 SLA Profile

The following syntax applies to the 7450 ESS:

CLI Syntax:

```
configure>service>ies service-id
interface ip-int-name
     sap sap-id
          host {[ip ip-address] [mac ieee-address}
             [subscriber sub-ident-string] [sub-profile
             sub-profile-name] [sla-profile sla-profile-
             name]
```

The following syntax applies to the 7750 SR:

CLI Syntax:
```
configure>service>ies service-id
interface ip-int-name
     sap sap-id
          host {[ip ip-address] [mac ieee-address}
             [subscriber sub-ident-string] [sub-profile
             sub-profile-name] [sla-profile sla-profile-
             name]
          sub-sla-mgmt
               def-sla-profile default-sla-profile-name
               single-sub-parameters
                    non-sub-traffic sub-profile sub-
                         profile-name sla-profile sla-
                         profile-name [subscriber sub-
                         ident-string]
subscriber-interface ip-int-name
          group-interface ip-int-name
               sap sap-id
                    host ip ip-address [mac ieee-address]
                         [subscriber sub-ident-string]
                         [sub-profile sub-profile-name]
                         [sla-profile sla-profile-name]
                    sub-sla-mgmt
                         def-sla-profile default-sla-
                           profile-name
                         single-sub-parameters
                              non-sub-traffic sub-
                                 profile sub-profile-name
                                 sla-profile sla-profile-
                                 name [subscriber sub-
                                 ident-string]
```

The following syntax applies to the 7450 ESS and 7750 SR:

CLI Syntax:
```
configure>service>vpls service-id
sap sap-id
     host {[ip ip-address] [mac ieee-address]}
        [subscriber sub-ident-string] [sub-profile sub-
        profile-name] [sla-profile sla-profile-name]
     sub-sla-mgmt
          def-sla-profile default-sla-profile-name
          single-sub-parameters
               non-sub-traffic sub-profile sub-profile-
                  name sla-profile sla-profile-name
                  [subscriber sub-ident-string]
```

The following syntax applies to the 7750 SR:

CLI Syntax:
```
configure>service>vprn service-id
interface ip-int-name
     sap sap-id
          host {[ip ip-address] [mac ieee-address]}
             [subscriber sub-ident-string] [sub-profile
             sub-profile-name] [sla-profile sla-profile-
             name]
```

The following syntax applies to the 7450 ESS and 7750 SR:

CLI Syntax:
```
configure>subscriber-mgmt
explicit-subscriber-map
     entry key sub-ident-string [sub-profile sub-
       profile-name] [alias sub-alias-string] [sla-
       profile sla-profile-name]
sub-ident-policy sub-ident-policy-name
     sla-profile-map
          entry key sla-profile-string sla-profile
            sla-profile-name
sub-profile sla-profile-map
     sla-profile-map
          entry key sla-profile-string sla-profile
            sla-profile
```

9.24.4 Configuring Dual Homing

The following displays an example of a dual homing configuration a. The configuration shows dual homing with a peer node with a system address of 10.1.1.23. The DHCP server returns a default route with a 10.21.1.3 next hop. This example only applies to the 7750 SR.

```
A:ALA-48#
#------------------------------------------------
echo "Redundancy Configuration"
#------------------------------------------------
    redundancy
        multi-chassis
            peer 10.1.1.23 create
                sync
                    srrp
                    sub-mgmt
                    port lag-100 sync-tag "Tag1" create
                    exit
                    no shutdown
                exit
                no shutdown
            exit
        exit
    exit
```

```
#------------------------------------------------
echo "Service Configuration"
#------------------------------------------------
    service
        customer 1 create
            description "Default customer"
        exit
        sdp 23 create
            far-end 10.1.1.23
            no shutdown
        exit
        ies 40 customer 1 create
            redundant-interface "r40-1" create
                address 10.1.1.1/31
                spoke-sdp 23:1 create
                exit
            exit
            subscriber-interface "s40-1" create
                address 10.21.1.1/16 gw-ip-address 10.21.1.3
                dhcp
                    gi-address 10.21.1.1
                exit
                group-interface "g40-1" create
                    dhcp
                        server 10.1.1.1
                        lease-populate 8000
                        no shutdown
                    exit
                    redundant-interface r40-1
                    remote-proxy-arp
                    sap lag-100:1 create
                        sub-sla-mgmt
                            def-sub-profile "subProf"
                            def-sla-profile "slaProf"
                            sub-ident-policy "subIdentPolicy"
                            multi-sub-sap
                            no shutdown
                        exit
                    exit
                    sap lag-100:4094 create
                    exit
                    srrp 1 create
                        message-path lag-100:4094
                        no shutdown
                    exit
                exit
            exit
            no shutdown
        exit
exit
...
-----------------------------------------------
A:ALA-48#
```

9.24.4.1 SHCV Policies

Under the group-interface, the **host-connectivity-verify** configuration was used as a reference timer for some event triggered SHCV while other used hard-coded values. The SHCV-policy and separated out every type of SHCV and allowed each type to have their individual configurable timer values. Furthermore, individual SHCV trigger types can be shut down. The SHCV policy can be applied to one or more group interfaces and can be configured differently for IPv4 vs. IPv6 hosts.

There are various types of triggered SHCV:

- **ip-conflict**: Sent when a SAP detects that there is a IP address or prefix conflict on the SAP.
- **host-limit-exceeded**: Sent when a subscriber has exceeded a configured host-limit. Host-limits are set on sla-profile host-limit, ipoe-session sap-session-limit, and ipoe-session session-limit.
- **inactivity**: Category-map configured under sla-profile can trigger an SHCV once the subscriber host has become idle.
- **mobility**: Intended for mobility applications such as WiFi. When a subscriber moves between SAPs and requests for the same IP address, a triggered SHCV is sent to verify if the old host is still connected before removing the old host entry.
- **mac-learning**: For IP-only static-host MAC learning. The trigger SHCV is sent to learn the subscriber MAC when a **no shutdown** command is executed on the CLI for the static host.

Note that some SHCVs are triggered based on a host's DHCP messages. These DHCP messages are not buffered. The SHCV is used only to perform a verification check on an old host to verify if the host is still connected to the BNG. Therefore, the BNG still requires the new hosts to retransmit their DHCP messages once the SHCV removes the disconnected host.

9.24.4.2 SHCV Policy

Under the group-interface, the **host-connectivity-verify** configuration was used as a reference timer for some event triggered SHCV while other used hard-coded values. The SHCV-policy and separated out every type of SHCV and allowed each type to have their individual configurable timer values. Furthermore, individual SHCV trigger types can be shut down. The SHCV policy can be applied to one or more group interfaces and can be configured differently for IPv4 vs. IPv6 hosts.

There are various types of triggered SHCV:

- **ip-conflict**: Sent when a SAP detects that there is a IP address or prefix conflict on the SAP.
- **host-limit-exceeded**: Sent when a subscriber has exceeded a configured host limit. Host limits are set with **sla-profile>host-limits**, **ipoe-session>sap-session-limit**, and **ipoe-session>session-limit**.
- **inactivity**: **sla-profile>category-map** can trigger an SHCV once the subscriber host has become idle.
- **mobility**: Intended for mobility applications such as WiFi. When a subscriber moves between SAPs and requests for the same IP address, a triggered SHCV is sent to verify if the old host is still connected before removing the old host entry.
- **mac-learning**: For IP-only static-host MAC learning. The trigger SHCV is sent to learn the subscriber MAC when a **no shutdown** command is executed on the CLI for the static host.

Some SHCVs are triggered based on a host's DHCP messages. These DHCP messages are not buffered. The SHCV is used only to perform a verification check on an old host to verify if the host is still connected to the BNG. Therefore, the BNG still requires the new hosts to retransmit their DHCP messages once the SHCV removes the disconnected host.

9.24.4.3 Subscriber Identification Policy

The following syntax applies to the 7450 ESS:

CLI Syntax:
```
configure>service>ies service-id
interface ip-int-name
     sap sap-id
          host {[ip ip-address] [mac ieee-address}
             [subscriber sub-ident-string] [sub-profile
             sub-profile-name] [sla-profile sla-profile-
             name]
```

The following syntax applies to the 7750 SR:

CLI Syntax:
```
configure>service>ies service-id
interface ip-int-name
     sap sap-id
          host {[ip ip-address] [mac ieee-address}
             [subscriber sub-ident-string] [sub-profile
             sub-profile-name] [sla-profile sla-profile-
             name]
          sub-sla-mgmt
               single-sub-parameters
```

```
                    non-sub-traffic sub-profile sub-
                         profile-name sla-profile sla-
                         profile-name [subscriber sub-
                         ident-string]
               sub-ident-policy sub-ident-policy-name
```

9.24.4.4 Subscriber Profile

The following syntax applies to the 7450 ESS:

CLI Syntax:
```
configure>service>ies service-id
interface ip-int-name
     sap sap-id
          host {[ip ip-address] [mac ieee-address}
             [subscriber sub-ident-string] [sub-profile
             sub-profile-name] [sla-profile sla-profile-
             name]
```

The following syntax applies to the 7750 SR:

CLI Syntax:
```
configure>service>ies service-id
interface ip-int-name
     sap sap-id
          host {[ip ip-address] [mac ieee-address}
             [subscriber sub-ident-string] [sub-profile
             sub-profile-name] [sla-profile sla-profile-
             name]
          sub-sla-mgmt
               def-sub-profile default-subscriber-
                  profile-name
               single-sub-parameters
                    non-sub-traffic sub-profile sub-
                         profile-name sla-profile sla-
                         profile-name [subscriber sub-
                         ident-string]
```

The following syntax applies to the 7450 ESS and 7750 SR:

CLI Syntax:
```
configure>service>vpls service-id
sap sap-id
     host {[ip ip-address] [mac ieee-address]}
        [subscriber sub-ident-string] [sub-profile sub-
        profile-name] [sla-profile sla-profile-name]
     sub-sla-mgmt
          def-sub-profile default-sub-profile-name
          single-sub-parameters
```

```
non-sub-traffic sub-profile sub-profile-
  name sla-profile sla-profile-name
  [subscriber sub-ident-string]
```

The following syntax applies to the 7450 ESS and 7750 SR:

CLI Syntax:

```
configure>subscriber-mgmt
sub-profile subscriber-profile-name
explicit-subscriber-map
     entry key sub-ident-string [sub-profile sub-
       profile-name] [alias sub-alias-string] [sla-
       profile sla-profile-name]
sub-ident-policy sub-ident-policy-name
     sub-profile-map
          entry key sub-profile-string sub-profile sub-
            profile-name
```

9.25 Subscriber Management Command Reference

9.25.1 Command Hierarchies

- ANCP Commands
- Authentication Policy Commands
- BGP Peering Policy Commands
- Call Trace Commands
- Category Map and Credit Control Policy Commands
- Diameter Policy Commands
- Explicit Subscriber Mapping Commands
- Filter Commands
- GSMP Configuration Commands
- Host Lockout Commands
- Host Tracking Policy Commands
- IGMP Policy Commands
- Managed SAP Policy Commands
- Multicast Listener Discovery (MLD) Commands
- Multi-Chassis Redundancy Commands
- PIM Policy Commands
- RADIUS Accounting Policy Commands
- RADIUS Route Download Commands
- Redundant Interface Commands
- RIP Commands
- Router Advertisement Commands
- SHCV Policy Commands
- SLA Profile Commands
- Subscriber Identification Policy Commands
  - Auto-Generated Subscriber Identification Key Commands
  - Auto-Generated Subscriber Identification Key Service Commands
- Subscriber Management Diameter Application Policy Commands
- Subscriber Management Service Commands

- IES Subscriber Management Configuration Commands
  - IES Group Interface Commands
- VPLS Subscriber Management Configuration Commands
- VPRN Subscriber Interface Commands
  - VPRN Group Interface Commands
- Subscriber MCAC Policy Commands
- Subscriber Profile Commands
- VPort Commands
- Wireless Portal Protocol (WPP) Commands
- Multiple PPPoE Session QoS Commands
- Show Commands
- Monitor Commands
- Clear Commands
- Debug Commands
- Tools Commands

Note:

- Enhanced Subscriber Management is supported on the redundant chassis model only.
- Refer to the Triple Play Service Commands section for information about service feature descriptions, CLI syntax, and usage.
- Refer to the DHCP Management chapter for information about DHCP feature descriptions, CLI syntax, and usage.

9.25.1.1 ANCP Commands

```
config
    — subscriber-mgmt
        — ancp
            — ancp-policy policy-name
                — egress
                    — rate-adjustment adjusted-percent
                    — no rate-adjustment
                    — rate-modify agg-rate-limit
                    — rate-modify scheduler scheduler-name
                    — no rate-modify
                    — rate-monitor kilobit-per-second [alarm]
                    — no rate-monitor
                    — rate-reduction kilobit-per-second
                    — no rate-reduction
                — ingress
```

```
                        — rate-adjustment adjusted-percent
                        — no rate-adjustment
                        — rate-modify scheduler scheduler-name
                        — no rate-modify
                        — rate-monitor kilobit-per-second [alarm]
                        — no rate-monitor
                        — rate-reduction kilobit-per-second
                        — no rate-reduction
                    — [no] port-down
                        — [no] disable-shcv [alarm] [hold-time seconds]
            — ancp-static-map
                - entry key ancp-string customer customer-id multi-service-site
                    customer-site-name ancp-policy policy-name
                - entry key ancp-string sap sap-id ancp-policy policy-name
                - no entry key ancp-string customer customer-id multi-service-site
                    customer-site-name
                - no entry key ancp-string sap sap-id
```

9.25.1.2 Authentication Policy Commands

```
config
    — aaa
        — radius-coa-port {port-number}
        — no radius-coa-port
    — subscriber-mgmt
        — [no] authentication-policy name
            — [no] accept-authorization-change
            — accept-script-policy policy-name
            — no accept-script-policy
            — coa-script-policy policy-name
            — no coa-script-policy
            — description description-string
            — no description
            — fallback-action accept [force-probing]
            — fallback-action user-db local-user-name [force-probing]
            — no fallback-action
            — gtp-user-name {imsi | imsi-apn}
            — no gtp-user-name
            — [no] include-radius-attribute
                — [no] access-loop-options
                — [no] acct-session-id
                — [no] apn
                — [no] called-station-id
                — calling-station-id
                — calling-station-id {llid | mac | remote-id | sap-id | sap-string}
                — no calling-station-id
                — [no] circuit-id
                — [no] dhcp-options
                — [no] dhcp-vendor-class-id
                — [no] dhcp6-options
                — [no] gprs-negotiated-qos-profile
```

- [no] imei
- [no] mac-address
- [no] msisdn
- [no] nas-identifier
- nas-port-id [prefix-string *string*] [suffix *suffix-option*]
- nas-port-type
- nas-port-type [*value*]
- no nas-port-type
- [no] pppoe-service-name
- [no] rat-type
- [no] remote-id
- [no] sap-session-index
- [no] tunnel-server-attrs
- [no] uli
- [no] wifi-num-attached-ues
- [no] wifi-ssid-vlan

- password *password*
- no password
- ppp-user-name append *domain-name*
- ppp-user-name default-domain *domain-name*
- ppp-user-name replace *domain-name*
- ppp-user-name strip
- no ppp-user-name
- pppoe-access-method {none | padi | pap-chap}
- no pppoe-access-method
- radius-accounting-server
  - access-algorithm {direct | round-robin}
  - no access-algorithm
  - [no] access-loop-options
  - hold-down-time *seconds*
  - no hold-down-time
  - retry *count*
  - no retry
  - router *router-instance*
  - router service-name *service-name*
  - no router
  - server *server-index* address *ip-address* secret key [hash | hash2] [port *port-num*] [coa-only] [pending-requests-limit *limit*]
  - no server *server-index*
  - source-address *ip-address*
  - no source-address
  - timeout *seconds*
  - no timeout
- radius-server-policy *radius-server-policy-name*
- [no] radius-server-policy
- [no] re-authentication
- request-script-policy *policy-name*
- no request-script-policy
- send-acct-stop-on-fail {[on-request-failure] [on-reject] [on-accept-failure]}
- no send-acct-stop-on-fail
- user-name-format *format* [mac-format *mac-format*]
- user-name-format *format* append [*domain-name*] [mac-format *mac-format*]
- user-name-format *format* append domain-name

— **user-name-format** *format* **default-domain** *domain-name* [**mac-format** *mac-format*]
— **user-name-format** *format* **replace** *domain-name* [**mac-format** *mac-format*]
— **user-name-format** *format* **strip** [**mac-format** *mac-format*]
— **no user-name-format**

9.25.1.3 BGP Peering Policy Commands

Note: The commands listed in this section apply only to the 7750 SR.

config
— **subscriber-mgmt**
— **bgp-peering-policy** *policy-name* [**create**]
— **no bgp-peering-policy** *policy-name*
— [**no**] **advertise-inactive**
— [**no**] **aggregator-id-zero**
— [**no**] **as-override**
— **auth-keychain** *name*
— **no auth-keychain**
— **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
— **no authentication-key**
— **cluster** *cluster-id*
— **no cluster**
— [**no**] **connect-retry** *seconds*
— [**no**] **damping**
— **description** *description-string*
— **no description**
— [**no**] **disable-4byte-asn**
— [**no**] **disable-client-reflect**
— **disable-communities** [**standard**] [**extended**]
— **no disable-communities**
— [**no**] **disable-fast-external-failover**
— **export** *policy-name* [*policy-name*]
— **no export**
— **hold-time** *seconds*
— **no hold-time**
— **import** *policy-name* [*policy-name*]
— **no import**
— **keepalive** *seconds*
— **no keepalive**
— **local-address** *ip-address*
— **no local-address**
— **local-as** *as-number* [**private**]
— **no local-as**
— **local-preference** *local-preference*
— **no local-preference**
— **loop-detect** {**drop-peer** | **discard-route** | **ignore-loop**| **off**}

— **no loop-detect**
— **med-out** {*number* | **igp-cost**}
— **no med-out**
— **min-route-advertisement** *seconds*
— **no min-route-advertisement**
— **multihop** *ttl-value*
— **no multihop**
— [**no**] **next-hop-self**
— [**no**] **passive**
— **peer-as** *as-number*
— **preference** *preference*
— **no preference**
— **prefix-limit** *limit*
— **no prefix-limit**
— [**no**] **remove-private**
— **ttl-security** *min-ttl-value*
— **no ttl-security**
— **type** {**internal** | **external**}
— **no type**

9.25.1.4 Call Trace Commands

configure
— **call-trace**
— **location** {**cf1** | **cf2**}
— [**no**] **disable**
— **size-limit** *limit-value*
— **no size-limit**
— **max-files-number** *number*
— **primary-cf** *cflash-id*
— **trace-profile** *profile-name* [**create**]
— **no trace-profile** *profile-name*
— **applications** [**connectivity-management**] [**radius-auto**] [**radius-acct**] [**python**] [**ludb**] [**msap**]
— **applications all**
— [**no**] **debug-output**
— **description** *description-string*
— **no description**
— **events** {**none** | **public-only** | **all**}
— **live-output** {*ip-address* | *fqdn*} [**port** *port*] [**router** {*router-instance* | **service-name** *service-name*}]
— **no live-output**
— **size-limit** *limit-value*
— **time-limit** *limit-value*

9.25.1.5 Category Map and Credit Control Policy Commands

config

```
— subscriber-mgmt
    — category-map category-map-name [create]
    — no category-map category-map-name
        — description description-string
        — no description
        — activity-threshold kilobits-per-second
        — no activity-threshold
        — category category-name [create]
        — no category category-name
            — credit-type-override {volume | time}
            — no credit-type-override
            — default-credit time seconds
            — default-credit volume credits [bytes | kilobytes | megabytes |
                gigabytes]
            — no default-credit
            — description description-string
            — no description
            — [no] exhausted-credit-service-level
                — [no] egress-ip-filter-entries
                    — entry entry-id [create]
                        — action drop
                        — action forward
                        — action http-redirect url [allow-override]
                        — no action
                        — description description-string
                        — no description
                        — match [protocol protocol-id]
                        — no match
                            — dscp dscp-name
                            — no dscp
                            — dst-port {lt | gt | eq} dst-port-number
                            — dst-port range start end
                            — no dst-port
                            — fragment {true | false}
                            — icmp-code icmp-code
                            — no icmp-code
                            — icmp-type icmp-type
                            — no icmp-type
                            — ip-option ip-option-value [ip-option-mask]
                            — no ip-option
                            — multiple-option {true | false}
                            — option-present {true | false}
                            — src-ip {ip-address/mask | ip-address
                                netmask}
                            — no src-ip
                            — src-port {lt | gt | eq} src-port-number
                            — src-port range start end
                            — no src-port
                            — tcp-ack {true | false}
                            — no tcp-ack
                            — tcp-syn {true | false}
                            — no tcp-syn
                — [no] egress-ipv6-filter-entries
                    — entry entry-id [create]
```

```
                    — action drop
                    — action forward
                    — no action
                    — description description-string
                    — no description
                    — match [next-header next-header]
                    — no match
                        — dscp dscp-name
                        — no dscp
                        — dst-port {lt | gt | eq} dst-port-number
                        — dst-port range start end
                        — no dst-port
                        — icmp-code icmp-code
                        — no icmp-code
                        — icmp-type icmp-type
                        — no icmp-type
                        — src-ip {ipv6-address | prefix-length}
                        — no src-ip
                        — src-port {lt | gt | eq} src-port-number
                        — src-port range start end
                        — no src-port
                        — tcp-ack {true | false}
                        — no tcp-ack
                        — tcp-syn {true | false}
                        — no tcp-syn
— [no] ingress-ip-filter-entries
    — entry entry-id [create]
        — action drop
        — action forward
        — action http-redirect url [allow-override]
        — no action
        — description description-string
        — no description
        — match [protocol protocol-id]
        — no match
            — dscp dscp-name
            — no dscp
            — dst-ip {ip-address/mask | ip-address
                netmask}
            — no dst-ip
            — dst-port {lt | gt | eq} dst-port-number
            — dst-port range start end
            — no dst-port
            — fragment {true | false}
            — icmp-code icmp-code
            — no icmp-code
            — icmp-type icmp-type
            — no icmp-type
            — ip-option ip-option-value [ip-option-mask]
            — no ip-option
            — multiple-option {true | false}
            — option-present {true | false}
            — src-port {lt | gt | eq} src-port-number
            — src-port range start end
```

```
                                        — no src-port
                                        — tcp-ack {true | false}
                                        — no tcp-ack
                                        — tcp-syn {true | false}
                                        — no tcp-syn
                                — [no] ingress-ipv6-filter-entries
                                    — entry entry-id [create]
                                        — action drop
                                        — action forward
                                        — no action
                                        — description description-string
                                        — no description
                                        — match [next-header next-header]
                                        — no match
                                            — dscp dscp-name
                                            — no dscp
                                            — dst-ip {ipv6-address | prefix-length}
                                            — no dst-ip
                                            — dst-port {lt | gt | eq} dst-port-number
                                            — dst-port range start end
                                            — no dst-port
                                            — icmp-code icmp-code
                                            — no icmp-code
                                            — icmp-type icmp-type
                                            — no icmp-type
                                            — src-port {lt | gt | eq} src-port-number
                                            — src-port range start end
                                            — no src-port
                                            — tcp-ack {true | false}
                                            — no tcp-ack
                                            — tcp-syn {true | false}
                                            — no tcp-syn
                                        — pir pir-rate
                                        — pir max
                                        — no pir
                                    — out-of-credit-action-override {continue | block-category | change-
                                      service-level}
                                    — no out-of-credit-action-override
                                    — policer policer-id {ingress-only | egress-only | ingress-egress}
                                    — no policer policer-id
                                    — queue queue-id {ingress-only | egress-only | ingress-egress}
                                    — no queue queue-id
                                    — rating-group rating-group-id
                                    — no rating-group
                                — credit-exhaust-threshold threshold-percentage
                                — no credit-exhaust-threshold
                                — credit-type {volume | time}
                                — no credit-type
                                — [no] gx-session-level-usage
                        — credit-control-policy policy-name [create]
                        — no credit-control-policy policy-name
                            — credit-control-server radius
                            — credit-control-server diameter policy-name
                            — no credit-control-server
```

— **default-category-map** *category-map-name*
— **no default-category-map**
— **description** *description-string*
— **no description**
— **error-handling-action** {**continue** | **block**}
— **no error-handling-action**
— **out-of-credit-action** *action*
— **no out-of-credit-action**

9.25.1.6 Diameter Policy Commands

- AAA Diameter Peer Policy Commands
- Subscriber Management Diameter Application Policy Commands
- Diameter Base Commands

9.25.1.6.1 AAA Diameter Peer Policy Commands

config
— **aaa**
— **diameter-peer-policy** *peer-policy-name* [**role** {**client** | **proxy**}] [**create**]
— **no diameter-peer-policy** *peer-policy-name*
— **applications** {[**gx**] [**gy**] [**nasreq**]}
— **no applications**
— **connection-timer** *connection-time*
— **no connection-timer**
— **description** *description-string*
— **no description**
— **ipv6-source-address** *ipv6-address*
— **no ipv6-source-address**
— **origin-host** *origin-host-string*
— **no origin-host**
— **origin-realm** *origin-realm-string*
— **no origin-realm**
— **peer** *name* [**create**]
— **no peer** *name*
— **address** *ip-address*
— **no address**
— **connection-timer** *connection-time*
— **no connection-timer**
— **destination-host** *destination-host-string*
— **no destination-host**
— **destination-realm** *destination-realm-string*
— **no destination-realm**
— **preference** *preference*
— **no preference**
— [**no**] **shutdown**
— **transaction-timer** *seconds*

```
            — no transaction-timer
            — transport tcp port port
            — no transport
            — watchdog-timer seconds
            — no watchdog-timer
        — proxy
            — mcs-peer ip-address sync-tag name
            — no mcs-peer
            — router router-instance address ip-address
            — no router
            — [no] shutdown
        — python-policy [python-policy-name]
        — no python-policy
        — router service service-name
        — router router-instance
        — no router
        — source-address ip-address
        — no source-address
        — transaction-timer transaction-time
        — no transaction-timer
        — vendor-support {three-gpp | vodafone}
        — no vendor-support
        — watchdog-timer wd-time
        — no watchdog-timer
```

9.25.1.6.2 Subscriber Management Diameter Application Policy Commands

```
configure
    — subscriber-mgmt
        — diameter-application-policy application-policy-name [create]
        — no diameter-application-policy application-policy-name
            — application {gx | gy | nasreq}
            — no application
            — description description-string
            — no description
            — diameter-peer-policy peer-policy-name
            — no diameter-peer-policy
            — gx
                — 3gpp-qos-mapping
                    — apn-ambr-dl agg-rate
                    — apn-ambr-dl arbiter arbiter-name
                    — apn-ambr-dl hs-sla-agg-rate
                    — apn-ambr-dl policer policer-id
                    — apn-ambr-dl queue queue-id
                    — apn-ambr-dl scheduler scheduler-name
                    — no apn-ambr-dl
                    — apn-ambr-ul arbiter arbiter-name
                    — apn-ambr-ul policer policer-id
                    — apn-ambr-ul queue queue-id
                    — apn-ambr-ul scheduler scheduler-name
                    — no apn-ambr-ul
```

- **avp-subscription-id** *origin* [**type** *type*]
- **no avp-subscription-id**
- **ccrt-replay**
  - **interval** *seconds*
  - **no interval**
  - **max-lifetime** *hours*
  - **no max-lifetime**
  - [**no**] **shutdown**
- [**no**] **include-avp**
  - [**no**] **an-gw-address**
  - [**no**] **apn-ambr**
  - [**no**] **called-station-id**
  - **calling-station-id** [**type** {**llid** | **mac** | **remote-id** | **sap-id** | **sap-string**}]
  - **no calling-station-id**
  - [**no**] **ip-can-type**
  - [**no**] **logical-access-id**
  - **nas-port** *binary-spec*
  - **no nas-port**
  - **nas-port-id** [**prefix-type** {**none** | **user-string**}] [**prefix-string** *prefix-string*] [**suffix-type** {**circuit-id** | **none** | **remote-id** | **user-string**}] [**suffix-string** *suffix-string*]
  - **no nas-port-id**
  - **nas-port-type** [*type*]
  - **no nas-port-type**
  - [**no**] **pdn-connection-id**
  - [**no**] **physical-access-id**
  - [**no**] **rai**
  - [**no**] **rat-type**
  - [**no**] **sgsn-mcc-mnc**
  - [**no**] **supported-features**
  - **user-equipment-info** [**type** *ue-info-type*]
  - **no user-equipment-info**
  - [**no**] **user-location-info**
- **mac-format** *mac-format*
- **no mac-format**
- [**no**] **report-ip-address-event**

gy

- **avp-subscription-id** *origin* [**type** *type*]
- **no avp-subscription-id**
- **ccrt-replay**
  - **interval** *seconds*
  - **no interval**
  - **max-lifetime** *hours*
  - **no max-lifetime**
  - [**no**] **shutdown**
- **extended-failure-handling**
  - **interim-credit**
    - **max-attempts** *count*
    - **max-attempts infinite**
    - **no max-attempts**
    - [**no**] **reporting**
    - **validity-time** *seconds*
    - **no validity-time**

```
                                    — volume credits {bytes | kilobytes | megabytes |
                                        gigabytes}
                                    — no volume
                                — [no] new-session-id
                                — [no] shutdown
                            — mac-format mac-format
                            — no mac-format
                            — out-of-credit-reporting {final | quota-exhausted}
                            — no out-of-credit-reporting
                            — vendor-support {three-gpp | vodafone}
                            — no vendor-support
                            — [no] include-avp
                                — [no] 3gpp-charging-characteristics
                                — 3gpp-charging-id {auto | esm-info | id}
                                — no 3gpp-charging-id
                                — [no] 3gpp-ggsn-address
                                — [no] 3gpp-ggsn-ipv6-address
                                — [no] 3gpp-gprs-negotiated-qos-profile
                                — 3gpp-imsi {circuit-id | imsi | subscriber-id}
                                — no 3gpp-imsi
                                — [no] 3gpp-nsapi
                                — 3gpp-rat-type value
                                — no 3gpp-rat-type
                                — [no] 3gpp-selection-mode
                                — [no] 3gpp-session-stop-indicator
                                — [no] 3gpp-user-location-info
                                — [no] address-avp
                                — called-station-id [64 chars max]
                                — no called-station-id
                                — charging-rule-base-name category-map-name
                                — charging-rule-base-name string
                                — no charging-rule-base-name
                                — ggsn-address {ipv4 | ipv6}
                                — no ggsn-address
                                — [no] pdp-context-type
                                — [no] ps-information
                                — [no] radius-user-name
                                — service-context-id name
                                — no service-context-id
                                — user-equipment-info [type ue-info-type]
                                — no user-equipment-info
                    — nasreq
                        — [no] include-avp
                            — [no] called-station-id
                            — calling-station-id [type {llid | mac | remote-id | sap-id | sap-
                                string}]
                            — no calling-station-id
                            — [no] circuit-id
                            — [no] imei
                            — nas-port binary-spec
                            — no nas-port
                            — nas-port-id [prefix-type {none | user-string}] [prefix-string [8
                                chars max]] [suffix-type {circuit-id | none | remote-id |
                                user-string}] [suffix-string [64 chars max]]
```

```
                    — no nas-port-id
                    — nas-port-type [type]
                    — nas-port-type
                    — no nas-port-type
                    — [no] rat-type
                    — [no] remote-id
                    — [no] user-location-info
                — mac-format mac-format
                — no mac-format
                — password password
                — no password
                — user-name-format format
                — no user-name-format
                — user-name-operation operation [domain domain-name]
                — no user-name-operation
            — on-failure [failover {enabled | disabled}] [handling {continue | retry-and-
              terminate | terminate}]
            — no on-failure
            — on-failure (legacy) [failover {enabled | disabled}] [handling {continue |
              retry-and-terminate | terminate}]
            — no on-failure (legacy)
            — tx-timer seconds
            — no tx-timer
```

9.25.1.6.3 Diameter Base Commands

```
config
    — aaa
        — diameter
            — node origin-host-string [origin-realm origin-realm-string] [create]
            — no node origin-host-string
                    — connection-timer seconds
                    — no connection-timer
                    — description description-string
                    — no description
                    — ipv6-source-address ipv6-address
                    — no ipv6-source-address
                    — peer index index [destination-host-string] [create]
                    — no peer index index
                        — address [ip-address | ipv6-address]
                        — no address
                        — connection-timer seconds
                        — no connection-timer
                        — preference preference
                        — no preference
                        — [no] shutdown
                        — watchdog-timer seconds
                        — no watchdog-timer
                    — python-policy python-policy-name
                    — no python-policy
                    — router router-name
```

```
                        — router service service-name
                        — no router
                        — source-address ip-address
                        — no source-address
                — diameter-peer-policy peer-policy-name [role {client | proxy}] [create]
                — no diameter-peer-policy peer-policy-name

config
    — subscriber-mgmt
        — diameter-application-policy application-policy-name [create]
        — no diameter-application-policy application-policy-name
            — diameter-node origin-host-string destination-realm destination-realm-string
            — no diameter-node
            — diameter-peer-policy peer-policy-name
            — no diameter-peer-policy
            — on-failure (legacy) [failover {enabled | disabled}] [handling {continue |
                retry-and-terminate | terminate}]
            — no on-failure (legacy)

config
    — redundancy
        — multi-chassis
            — peer ip-address [create]
            — no peer ip-address
                — [no] sync
                    — [no] diameter-proxy
```

9.25.1.7 Explicit Subscriber Mapping Commands

```
config
    — subscriber-mgmt
        — explicit-subscriber-map
            — entry key sub-ident-string [sub-profile sub-profile-name] [alias sub-
                alias-string] [sla-profile sla-profile-name] [app-profile
                name]
            — no entry key sub-ident-string
```

9.25.1.8 Filter Commands

Refer to the 7450 ESS, 7750 SR, 7950 XRS, and VSR Router Configuration Guide for a complete list of IP filter and IPv6 filter commands and descriptions.

9.25.1.9 GSMP Configuration Commands

Note: The commands listed in this section apply only to the 7750 SR.

```
config
    — service
        — vpls service-id [customer customer-id] [create] [vpn vpn-id] [m-vpls] [b-vpls | i-vpls]
        — no vpls service-id
        — [no] vprn service-id [customer customer-id]  [vpn vpn-id] [m-vpls]
            — gsmp
                — [no] group name [create]
                    — ancp
                        — [no] dynamic-topology-discover
                        — [no] oam
                    — description description-string
                    — no description
                    — hold-multiplier multiplier
                    — no hold-multiplier
                    — [no] idle-filter
                    — keepalive seconds
                    — no keepalive
                    — [no] neighbor ip-address
                        — description description-string
                        — no description
                        — local-address ip-address
                        — no local-address
                        — priority-marking dscp dscp-name
                        — priority-marking prec ip-prec-value
                        — no priority-marking
                        — [no] shutdown
                    — [no] persistency-database
                    — [no] shutdown
                — [no] shutdown
```

9.25.1.10 Host Lockout Commands

Note: The commands listed in this section apply only to the 7750 SR.

```
config
    — subscriber-mgmt
        — host-lockout-policy policy-name [create]
        — no host-lockout-policy policy-name
```

```
            — description description-string
            — no description
            — host-key {mac}
            — no host-key
            — lockout-reset-time seconds
            — no lockout-reset-time
            — lockout-time [min seconds] [max seconds]
            — no lockout-time
            — max-lockout-hosts hosts
            — no max-lockout-hosts

config
    — service
        — vprn service-id [customer customer-id] [name name] [create]
        — no vprn service-id
            — [no] interface ip-int-name
                — [no] sap sap-id
                    — host-lockout-policy policy-name
                    — no host-lockout-policy
            — subscriber-interface ip-int-name [fwd-service service-id fwd-subscriber-
              interface ip-int-name] [create] [wan-mode mode]
            — no subscriber-interface ip-int-name
                — group-interface ip-int-name [create] [type]
                — no group-interface ip-int-name
                    — sap sap-id [create]
                    — no sap sap-id
                        — host-lockout-policy policy-name
                        — no host-lockout-policy

config
    — service
        — ies service-id [customer customer-id] [vpn vpn-id] [name name] [create]
            — [no] interface ip-int-name
                — [no] sap sap-id
                    — host-lockout-policy policy-name
                    — host-lockout-policy
            — subscriber-interface ip-int-name [fwd-service service-id fwd-subscriber-
              interface ip-int-name] [create] [wan-mode mode]
            — no subscriber-interface ip-int-name
                — [no] sap sap-id
                    — host-lockout-policy policy-name
                    — no host-lockout-policy

config
    — service
        — [no] vpls service-id [customer customer-id] [vpn vpn-id] [m-vpls] [create] [name
          name]
            — sap sap-id [split-horizon-group group-name] [create] [capture-sap] [eth-ring
              ring-index]
            — no sap sap-id
                — host-lockout-policy policy-name
                — no host-lockout-policy
```

9.25.1.11 Host Tracking Policy Commands

```
config
    — subscriber-mgmt
        — host-tracking-policy policy-name [create]
        — no host-tracking-policy policy-name
            — description description-string
            — no description
            — egress-rate-modify agg-rate-limit
            — egress-rate-modify scheduler scheduler-name
            — no egress-rate-modify
```

9.25.1.12 IGMP Policy Commands

Note: The commands listed in this section apply only to the 7750 SR.

```
config
    — subscriber-mgmt
        — igmp-policy policy-name [create]
            — description description-string
            — no description
            — [no] disable-router-alert-check
            — egress-rate-modify [agg-rate-limit | scheduler scheduler-name]
            — [no] fast-leave
            — import policy-name
            — no import
            — max-num-groups max-num-groups
            — no max-num-groups
            — max-num-sources max-num-sources
            — no max-num-sources
            — max-num-grp-sources [max-num-grp-sources]
            — no max-num-grp-sources
            — [no] mcast-reporting
                — mcast-reporting-dest dest-name
                — no mcast-reporting-dest
                — opt-reporting-fields [host-mac] [pppoe-session-id] [svc-id] [sap-id]
                — no opt-reporting-fields
                — [no] shutdown
            — per-host-replication [uni-mac | mcast-mac]
            — no per-host-replication
            — query-interval seconds
            — no query-interval
            — query-last-member-interval seconds
            — no query-last-member-interval
            — query-response-interval seconds
            — no query-response-interval
```

```
— redirection-policy policy-name
— no redirection-policy
— static
    — [no] group ip-address
— version version
— no version
```

9.25.1.13 Multi-Chassis Redundancy Commands

```
config
    — redundancy
        — multi-chassis
            — peer ip-address [create]
            — no peer ip-address
                — authentication-key [authentication-key | hash-key] [hash | hash2]
                — no authentication-key
                — description description-string
                — no description
                — mc-ipsec
                    — [no] bfd-enable
                    — discovery-interval interval-1 [boot interval-2]
                    — no discovery-interval
                    — hold-on-neighbor-failure multiplier
                    — no hold-on-neighbor-failure
                    — keep-alive-interval time-interval
                    — no keep-alive-interval
                    — tunnel-group group-id [create]
                    — no tunnel-group group-id
                        — peer-group group-id
                        — no peer-group
                        — priority priority
                        — no priority
                        — [no] preempt
                        — [no] shutdown
                — [no] mc-lag
                    — hold-on-neighbor-failure multiplier
                    — no hold-on-neighbor-failure
                    — keep-alive-interval interval
                    — no keep-alive-interval
                    — lag lag-id lacp-key admin-key system-id system-id [remote-lag
                          lag-id] system-priority system-priority
                    — no lag lag-id
                    — [no] shutdown
                — mc-ring
                    — [no] ring sync-tag [create]
                        — in-band-control-path
                            — [no] debounce
                            — dst-ip ip-address
                            — no dst-ip
                            — interface ip-int-name
                            — no interface
```

```
            — no max-debounce-time
            — service-id service-id
            — no service-id
            — service-name svc-name
            — no service-name
        — [no] path-b
            — [no] range vlan-range
        — [no] path-excl
            — [no] range vlan-range
        — [no] ring-node ring-node-name
            — connectivity-verify
                — dst-ip ip-address
                — no dst-ip
                — interval interval
                — no interval
                — service-id service-id
                — no service-id
                — service-name svc-name
                — no service-name
                — [no] shutdown
                — src-ip ip-address
                — no src-ip
                — src-mac ieee-address
                — no src-mac
                — vlan [vlan-encap]
                — no vlan
        — [no] shutdown
    — [no] l3-ring sync-tag
        — in-band-control-path
            — [no] debounce
            — dst-ip ip-address
            — no dst-ip
            — interface ip-int-name
            — no interface
            — max-debounce-time max-debounce-time
            — no max-debounce-time
            — service-id service-id
            — no service-id
            — service-name svc-name
            — no service-name
        — [no] ring-node ring-node-name
            — connectivity-verify
                — dst-ip ip-address
                — no dst-ip
                — interval interval
                — no interval
                — service-id service-id
                — no service-id
                — service-name svc-name
                — no service-name
                — [no] shutdown
                — src-ip ip-address
                — no src-ip
                — src-mac ieee-address
```

— **no src-mac**
— **vlan** [*vlan-encap*]
— **no vlan**
— [**no**] **srrp-instance** *srrp-id*
— [**no**] **shutdown**
— [**no**] **shutdown**
— **source-address** *ip-address*
— **no source-address**
— [**no**] **sync** **— Refer to the** *7450 ESS, 7750 SR, 7950 XRS, and VSR Basic System Configuration Guide* **for a list of sync commands and descriptions.**

9.25.1.14 PIM Policy Commands

config
— **subscriber-mgmt**
— **pim-policy** *policy-name* [**create**]
— **no pim-policy** *policy-name*
— **description** *description-string*
— **no description**

9.25.1.15 RADIUS Accounting Policy Commands

config
— **subscriber-mgmt**
— **radius-accounting-policy** *name* [**create**]
— **no radius-accounting-policy** *name*
— [**no**] **custom-record**
— **override-counter** *override-counter-id* [**create**]
— **no override-counter** *override-counter-id*
— **e-counters** [**all**]
— **no e-counters**
— [**no**] **in-profile-octets-discarded-count**
— [**no**] **in-profile-octets-forwarded-count**
— [**no**] **in-profile-packets-discarded-count**
— [**no**] **in-profile-packets-forwarded-count**
— [**no**] **out-profile-octets-discarded-count**
— [**no**] **out-profile-octets-forwarded-count**
— [**no**] **out-profile-packets-discarded-count**
— [**no**] **out-profile-packets-forwarded-count**
— **i-counters** [**all**]
— **no i-counters**
— [**no**] **all-octets-offered-count**
— [**no**] **all-packets-offered-count**
— [**no**] **high-octets-discarded-count**
— [**no**] **high-packets-discarded-count**
— [**no**] **in-profile-octets-forwarded-count**
— [**no**] **in-profile-packets-forwarded-count**

- [no] low-octets-discarded-count
- [no] low-packets-discarded-count
- [no] out-profile-octets-forwarded-count
- [no] out-profile-packets-forwarded-count

- queue *queue-id* [create]
- noqueue *queue-id*
  - e-counters [all]
  - no e-counters
    - [no] in-profile-octets-discarded-count
    - [no] in-profile-octets-forwarded-count
    - [no] in-profile-packets-discarded-count
    - [no] in-profile-packets-forwarded-count
    - [no] out-profile-octets-discarded-count
    - [no] out-profile-octets-forwarded-count
    - [no] out-profile-packets-discarded-count
    - [no] out-profile-packets-forwarded-count
  - i-counters [all]
  - no i-counters
    - [no] all-octets-offered-count
    - [no] all-packets-offered-count
    - [no] high-octets-discarded-count
    - [no] high-octets-offered-count
    - [no] high-packets-discarded-count
    - [no] high-packets-offered-count
    - [no] in-profile-octets-forwarded-count
    - [no] in-profile-packets-forwarded-count
    - [no] low-octets-discarded-count
    - [no] low-octets-offered-count
    - [no] low-packets-discarded-count
    - [no] low-packets-offered-count
    - [no] out-profile-octets-forwarded-count
    - [no] out-profile-packets-forwarded-count
    - [no] uncoloured-octets-offered-count
    - [no] uncoloured-octets-offered-count
- ref-override-counter *ref-override-counter-id*
- ref-override-counter all
- no ref-override-counter
  - e-counters [all]
  - no e-counters
    - [no] in-profile-octets-discarded-count
    - [no] in-profile-octets-forwarded-count
    - [no] in-profile-packets-discarded-count
    - [no] in-profile-packets-forwarded-count
    - [no] out-profile-octets-discarded-count
    - [no] out-profile-octets-forwarded-count
    - [no] out-profile-packets-discarded-count
    - [no] out-profile-packets-forwarded-count
  - i-counters [all]
  - no i-counters
    - [no] all-octets-offered-count
    - [no] all-packets-offered-count
    - [no] high-octets-discarded-count
    - [no] high-packets-discarded-count
    - [no] in-profile-octets-forwarded-count

- [no] in-profile-packets-forwarded-count
- [no] low-octets-discarded-count
- [no] low-packets-discarded-count
- [no] out-profile-octets-forwarded-count
- [no] out-profile-packets-forwarded-count

- ref-queue *queue-id*
- ref-queue all
- no ref-queue
  - e-counters [all]
  - no e-counters
    - [no] in-profile-octets-discarded-count
    - [no] in-profile-octets-forwarded-count
    - [no] in-profile-packets-discarded-count
    - [no] in-profile-packets-forwarded-count
    - [no] out-profile-octets-discarded-count
    - [no] out-profile-octets-forwarded-count
    - [no] out-profile-packets-discarded-count
    - [no] out-profile-packets-forwarded-count
  - i-counters [all]
  - no i-counters
    - [no] all-octets-offered-count
    - [no] all-packets-offered-count
    - [no] high-octets-discarded-count
    - [no] high-octets-offered-count
    - [no] high-packets-discarded-count
    - [no] high-packets-offered-count
    - [no] in-profile-octets-forwarded-count
    - [no] in-profile-packets-forwarded-count
    - [no] low-octets-discarded-count
    - [no] low-packets-discarded-count
    - [no] low-octets-offered-count
    - [no] low-packets-offered-count
    - [no] out-profile-octets-forwarded-count
    - [no] out-profile-packets-forwarded-count
    - [no] uncoloured-octets-offered-count
    - [no] uncoloured-packets-offered-count
  - significant-change *delta*
  - no significant-change
- description *description-string*
- no description
- [no] delay-start-time
- [no] host-accounting [interim-update]
- include-radius-attribute
  - [no] access-loop-options
  - [no] acct-authentic
  - [no] acct-delay-time
  - [no] alc-acct-triggered-reason
  - [no] alc-error-code
  - [no] all-authorized-session-addresses
  - [no] apn
  - [no] bearer-fteid
  - [no] bonding-active-connections
  - [no] bonding-id
  - [no] called-station-id

- **calling-station-id**
- **calling-station-id** {**llid** | **mac** | **remote-id** | **sap-id** | **sap-string**}
- **no calling-station-id**
- [**no**] **circuit-id**
- [**no**] **delegated-ipv6-prefix**
- [**no**] **detailed-acct-attributes**
- [**no**] **dhcp-vendor-class-id**
- [**no**] **framed-interface-id**
- [**no**] **framed-ip-addr**
- [**no**] **framed-ip-netmask**
- [**no**] **framed-ipv6-prefix**
- [**no**] **framed-ipv6-route**
- [**no**] **framed-route**
- [**no**] **imei**
- [**no**] **ipv6-address**
- [**no**] **mac-address**
- [**no**] **nas-identifier**
- [**no**] **msisdn**
- **nas-port** *binary-spec*
- **no nas-port**
- **nas-port-id**
- **nas-port-id** [**prefix-string** *string*] [**suffix** *suffix-option*]
- **no nas-port-id**
- **nas-port-type** [*value*]
- **no nas-port-type**
- [**no**] **nat-port-range**
- [**no**] **remote-id**
- [**no**] **sla-profile**
- [**no**] **spi-sharing-id**
- [**no**] **std-acct-attributes**
- [**no**] **sub-profile**
- [**no**] **subscriber-id**
- [**no**] **tunnel-server-attrs**
- [**no**] **uli**
- [**no**] **user-name**
- [**no**] **v6-aggregate-stats**
- [**no**] **wifi-rssi**

- **queue-instance-accounting** [**interim-update**]
- **no queue-instance-accounting**
- **radius-accounting-server**
  - **access-algorithm** {**direct** | **round-robin**}
  - **no access-algorithm**
  - **retry** *count*
  - **no retry**
  - **router** *router-instance*
  - **no router**
  - **server** *server-index* [**address** *ip-address*] [**secret** *key*] [**port** *port*] [**pending-requests-limit** *limit*]
  - **no server** *server-index*
  - **source-address** *ip-address*
  - **no source-address**
  - **timeout** *seconds*
  - **no timeout**
- **radius-server-policy** *radius-server-policy-name*

- [no] radius-server-policy
- session-accounting [interim-update] [host-update]
- no session-accounting
- session-id-format {description | number}
- no session-id-format
- triggered-updates
  - [no] gtp-mobility
- update-interval *minutes*
- no update-interval
- update-interval-jitter absolute *seconds*
- [no] update-interval-jitter

9.25.1.16 RADIUS Route Download Commands

```
configure
    aaa
        route-downloader name [create]
        no route-downloader name
            base-user-name user-name
            no base-user-name
            default-metric metric
            no default-metric
            default-tag tag
            no default-tag
            description description-string
            no description
            download-interval minutes
            no download-interval
            max-routes routes
            no max-routes
            password password [hash | hash2]
            no password
            radius-server-policy policy-name
            no radius-server-policy
            retry-interval min minimum max maximum
            no retry-interval
            [no] shutdown
```

9.25.1.17 Router Advertisement Commands

```
config
    service
        ies service-id [customer customer-id] [create] [vpn vpn-id] [name name]
        no ies service-id
        vprn service-id [name name] [customer customer-id] [create]
        no vprn service-id
            subscriber-interface ip-int-name [fwd-service service-id fwd-subscriber-
                interface ip-int-name] [create] [wan-mode mode]
```

```
— no subscriber-interface ip-int-name
    — [no] ipv6
        — [no] router-advertisements
            — current-hop-limit limit
            — no current-hop-limit
            — [no] dns-options
                — [no] include-dns
                — rdnss-lifetime seconds
                — rdnss-lifetime infinite
                — no rdnss-lifetime
            — force-mcast {ip | mac}
            — no force-mcast
            — [no] managed-configuration
            — max-advertisement seconds
            — no max-advertisement
            — min-advertisement seconds
            — no min-advertisement
            — mtu bytes
            — no mtu
            — [no] other-stateful-configuration
            — [no] prefix-options
                — [no] autonomous
                — [no] on-link
                — preferred-lifetime seconds
                — preferred-lifetime infinite
                — no preferred-lifetime
                — valid-lifetime seconds
                — valid-lifetime infinite
                — no valid-lifetime
            — reachable-time milli-seconds
            — no reachable-time
            — retransmit-time milli-seconds
            — no retransmit-time
            — router-lifetime seconds
            — router-lifetime no-default-router
            — no router-lifetime
            — [no] shutdown
            — [no] router-solicit
                — inactivity-timer [days days] [hrs hours] [min
                    minutes] [sec seconds]
                — inactivity-timer infinite
                — no inactivity-timer
    — [no] group-interface ip-int-name
        — [no] ipv6
            — [no] router-advertisements
                — current-hop-limit limit
                — no current-hop-limit
                — [no] dns-options
                    — [no] include-dns
                    — rdnss-lifetime seconds
                    — rdnss-lifetime infinite
                    — no rdnss-lifetime
                — force-mcast {ip | mac}
                — no force-mcast
```

```
                        — [no] managed-configuration
                        — max-advertisement seconds
                        — no max-advertisement
                        — min-advertisement seconds
                        — no min-advertisement
                        — mtu bytes
                        — no mtu
                        — [no] other-stateful-configuration
                        — [no] prefix-options
                            — [no] autonomous
                            — [no] on-link
                            — preferred-lifetime seconds
                            — preferred-lifetime infinite
                            — no preferred-lifetime
                            — valid-lifetime seconds
                            — valid-lifetime infinite
                            — no valid-lifetime
                        — reachable-time milli-seconds
                        — no reachable-time
                        — retransmit-time milli-seconds
                        — no retransmit-time
                        — router-lifetime seconds
                        — router-lifetime no-default-router
                        — no router-lifetime
                        — [no] shutdown
                    — [no] router-solicit
                        — inactivity-timer [days days] [hrs hours] [min
                            minutes] [sec seconds]
                        — inactivity-timer infinite
                        — no inactivity-timer

config
    — subscriber-mgmt
        — router-advertisement-policy name [create]
        — no router-advertisement-policy name
            — current-hop-limit limit
            — no current-hop-limit
            — [no] dns-options
                — [no] include-dns
                — rdnss-lifetime seconds
                — rdnss-lifetime infinite
                — no rdnss-lifetime
            — force-mcast {ip | mac}
            — no force-mcast
            — [no] managed-configuration
            — max-advertisement seconds
            — no max-advertisement
            — min-advertisement seconds
            — no min-advertisement
            — mtu bytes
            — no mtu
            — [no] other-stateful-configuration
            — [no] prefix-options
                — [no] stateful
```

```
                    — [no] auto-lifetimes
                    — [no] on-link
                    — preferred-lifetime seconds
                    — preferred-lifetime infinite
                    — no preferred-lifetime
                    — valid-lifetime seconds
                    — valid-lifetime infinite
                    — no valid-lifetime
                — [no] stateless
                    — [no] on-link
                    — preferred-lifetime seconds
                    — preferred-lifetime infinite
                    — no preferred-lifetime
                    — valid-lifetime seconds
                    — valid-lifetime infinite
                    — no valid-lifetime
            — reachable-time milli-seconds
            — no reachable-time
            — retransmit-time milli-seconds
            — no retransmit-time
            — router-lifetime seconds
            — router-lifetime no-default-router
            — no router-lifetime
```

9.25.1.18 SLA Profile Commands

```
config
    — subscriber-mgmt
        — sla-profile sla-profile-name [create]
        — no sla-profile sla-profile-name
            — category-map category-map-name [create]
            — no category-map
                — category category-name [create]
                — no category category-name
                    — idle-timeout timeout
                    — no idle-timeout
                    — idle-timeout-action {shcv-check | terminate}
                    — no idle-timeout-action
            — def-instance-sharing {per-sap | per-session}
            — description description-string
            — no description
            — egress
                — bonding-selection
                    — rate-thresholds high high-percentage low low-percentage
                    — no rate-thresholds
                    — weight initial-weight-percentage weight-change-percentage
                    — no weight
                — hs-agg-rate-limit kilobits-per-second
                — no hs-agg-rate-limit
                — hs-low-burst-max-class class
                — no hs-low-burst-max-class
```

- [no] ip-filter *filter-id*
- qos *policy-id* [vport-scheduler | port-scheduler] [force]
- no qos
  - hs-queue-stat-mode *mode*
  - no hs-queue-stat-mode
  - [no] hs-wrr-group *group-id*
    - hs-class-weight *weight*
    - no hs-class-weight
    - rate *rate*
    - no rate
  - policer *policer-id*
  - no policer *policer-id*
    - cbs *size* [bytes | kilobytes]
    - no cbs
    - mbs *size* [bytes | kilobytes]
    - no mbs
    - packet-byte-offset {add *bytes* | subtract *bytes*}
    - no packet-byte-offset
    - rate {max | *rate*} [cir {max | *rate*}]
    - stat-mode *stat-mode*
    - no stat-mode
  - queue *queue-id*
  - no queue *queue-id*
    - avg-frame-overhead *percent*
    - no avg-frame-overhead
    - cbs *size-in-kbytes*
    - no cbs
    - high-prio-only *percent*
    - no high-prio-only
    - hs-class-weight *weight*
    - no hs-class-weight
    - hs-wred-queue-policy *name*
    - no hs-wred-queue-policy
    - hs-wrr-weight *weight*
    - no hs-wrr-weight
    - mbs *size* [bytes | kilobytes]
    - no mbs
    - rate *pir-rate* [cir *cir-rate*]
    - no rate
    - stat-mode *stat-mode*
    - no stat-mode
  - [no] qos-marking-from-sap
  - report-rate agg-rate-limit
  - report-rate scheduler *scheduler-name*
  - report-rate pppoe-actual-rate
  - report-rate policer *policer-id*
  - report-rate rfc5515-actual-rate
  - no report-rate
  - scheduler-policy *scheduler-name* rate *pir-rate* [cir *cir-rate*]
  - no scheduler-policy
  - [no] use-ingress-l2tp-dscp
- [no] host-limits
  - ipv4-arp *max-nr-of-hosts*
  - no ipv4-arp

```
        — ipv4-dhcp max-nr-of-hosts
        — no ipv4-dhcp
        — ipv4-overall max-nr-of-hosts
        — no ipv4-overall
        — ipv4-ppp max-nr-of-hosts
        — no ipv4-ppp
        — ipv6-overall max-nr-of-hosts
        — no ipv6-overall
        — ipv6-pd-ipoe-dhcp max-nr-of-hosts
        — no ipv6-pd-ipoe-dhcp
        — ipv6-pd-overall max-nr-of-hosts
        — no ipv6-pd-overall
        — ipv6-pd-ppp-dhcp max-nr-of-hosts
        — no ipv6-pd-ppp-dhcp
        — ipv6-wan-ipoe-dhcp max-nr-of-hosts
        — no ipv6-wan-ipoe-dhcp
        — ipv6-wan-ipoe-slaac max-nr-of-hosts
        — no ipv6-wan-ipoe-slaac
        — ipv6-wan-overall max-nr-of-hosts
        — no ipv6-wan-overall
        — ipv6-wan-ppp-dhcp max-nr-of-hosts
        — no ipv6-wan-ppp-dhcp
        — ipv6-wan-ppp-slaac max-nr-of-hosts
        — no ipv6-wan-ppp-slaac
        — lac-overall max-nr-of-hosts
        — no lac-overall
        — overall max-nr-of-hosts
        — no overall
        — [no] remove-oldest
    — ingress
        — [no] ip-filter filter-id
        — qos policy-id [shared-queuing | multipoint-shared | service-queuing]
            [force]
        — no qos
            — queue queue-id
            — no queue queue-id
                — cbs size [bytes | kilobytes]
                — no cbs
                — high-prio-only percent
                — no high-prio-only
                — mbs size [bytes | kilobytes]
                — no mbs
                — rate pir-rate [cir cir-rate]
                — no rate
                — stat-mode {v4-v6}
                — no stat-mode
            — policer policer-id
            — no policer policer-id
                — cbs {size [bytes | kilobytes] | default}
                — no cbs
                — mbs size [bytes | kilobytes]
                — no mbs
                — packet-byte-offset {add bytes | subtract bytes}
                — no packet-byte-offset
```

```
                        — rate {max | rate} [cir {max | rate}]
                        — no rate
                        — stat-mode stat-mode
                        — no stat-mode
                    — report-rate agg-rate-limit
                    — report-rate scheduler scheduler-name
                    — report-rate pppoe-actual-rate
                    — report-rate policer policer-id
                    — report-rate rfc5515-actual-rate
                    — no report-rate
                — one-time-http-redirection ip-filter filter-id
```

9.25.1.19 Subscriber Identification Policy Commands

```
config
    — subscriber-mgmt
        — [no] sub-ident-policy sub-ident-policy-name [create]
            — app-profile-map
                — entry key app-profile-string app-profile app-profile-name
                — no entry key app-profile-string
                — [no] use-direct-map-as-default
            — description description-string
            — no description
            — primary
                — script-url dhcp-primary-script-url
                — no script-url
                — [no] shutdown
            — secondary
                — script-url dhcp-secondary-script-url
                — no script-url
                — [no] shutdown
            — sla-profile-map
                — entry key sla-profile-string sla-profile sla-profile-name
                — no entry key sla-profile-string
                — [no] use-direct-map-as-default
            — sub-profile-map
                — entry key sub-profile-string sub-profile sub-profile-name
                — no entry key sub-profile-string
                — [no] use-direct-map-as-default
            — tertiary
                — script-url dhcp-tertiary-script-url
                — no script-url
                — [no] shutdown
```

9.25.1.20 Auto-Generated Subscriber Identification Key Commands

Note: These commands are supported on the 7450 ESS in mixed mode.

```
config
    — subscriber-mgmt
        — auto-sub-id-key
            — ipoe-sub-id-key sub-id-key [sub-id-key]
            — no ipoe-sub-id-key
            — ppp-sub-id-key sub-id-key [sub-id-key]
            — no ppp-sub-id-key
```

9.25.1.21 Auto-Generated Subscriber Identification Key Service Commands

Note: Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* and *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* for further services commands.

```
config
    — service
        — vprn
            — subscriber-interface ip-int-name
                — group-interface ip-int-name
                    — sap sap-id
                        — sub-sla-mgmt
                            — def-sub-id use-auto-id
                            — def-sub-id use-sap-id
                            — def-sub-id string sub-id
                            — no def-sub-id
        — vpls
            — sap sap-id
                — sub-sla-mgmt
                    — def-sub-id use-auto-id
                    — def-sub-id use-sap-id
                    — def-sub-id string sub-id
                    — no def-sub-id
        — ies
            — subscriber-interface ip-int-name
                — group-interface ip-int-name
                    — sap sap-id
                        — sub-sla-mgmt
                            — def-sub-id use-auto-id
```

```
                                        — def-sub-id use-sap-id
                                        — def-sub-id string sub-id
                                        — no def-sub-id
        — subscriber-management
            — msap-policy msap-policy-name
                — sub-sla-mgmt
                                — def-sub-id use-auto-id
                                — def-sub-id use-sap-id
                                — def-sub-id string sub-id
                                — no def-sub-id
```

9.25.1.22 Subscriber MCAC Policy Commands

```
config
    — subscriber-mgmt
        — sub-mcac-policy sub-mcac-policy-name [create]
        — no sub-mcac-policy sub-mcac-policy-name
            — description description-string
            — no description
            — [no] shutdown
            — unconstrained-bw bandwidth mandatory-bw mandatory-bw
            — no unconstrained-bw
```

9.25.1.23 Subscriber Profile Commands

For information about configuring accounting policies, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR System Management Guide*.

```
config
    — subscriber-mgmt
        — accu-stats-policy policy-name [create]
        — no accu-stats-policy policy-name
            — description description-string
            — no description
            — [no] entry direction direction type type id id
        — [no] sub-profile subscriber-profile-name [create]
            — no accounting-policy
            — accu-stats-policy policy-name
            — no accu-stats-policy
            — ancp
                — ancp-policy name
                — no ancp-policy
            — [no] collect-stats
            — description description-string
            — no description
            — egress
                — agg-rate-limit agg-rate [queue-frame-based-accounting]
                — no agg-rate-limit
```

- **avg-frame-size** *bytes*
- **no avg-frame-size**
- **encap-offset** [**type** *type*]
- **no encap-offset**
- **hs-agg-rate-limit** *kilobits-per-second*
- **no hs-agg-rate-limit**
- **lag-per-link-hash class** {*class*} **weight** *weight*
- **no lag-per-link-hash**
- **policer-control-policy** *policy-name*
- **no policer-control-policy**
  - **max-rate** {*rate* | **max**}
  - **no max-rate**
  - **priority-mbs-thresholds**
    - **min-thresh-separation** *size* [**bytes** | **kilobytes**]
    - **no min-thresh-separation**
    - **priority** *level*
      - **mbs-contribution** *size* [**bytes** | **kilobytes**] [**fixed**]
- **scheduler-policy** *scheduler-policy-name*
- **no scheduler-policy**
  - **scheduler** *scheduler-name* **rate** [*pir-rate*] [**cir** *cir-rate*]
  - **no scheduler** *scheduler-name*

firewall-policy *policy-name*
no firewall-policy
host-tracking-policy *policy-name*
no host-tracking-policy
hs-sla-mode {**expanded** | **single**}
no hs-sla-mode
hsmda

- **egress-qos**
  - **agg-rate-limit** *agg-rate*
  - **no agg-rate-limit**
  - **qos** *policy-id*
  - **no qos**
    - **wrr-policy** *weight*
    - **no wrr-policy**
    - **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}
    - **no packet-byte-offset**
    - **queue** *queue-id*
    - **no queue** *queue-id*
      - **mbs** *size* [**bytes** | **kilobytes**]
      - **no mbs**
      - **rate** *pir-rate*
      - **no rate**
      - **rate** {**max** | *kilobits-per-second*}
      - **no rate**
      - **slope-policy** *hsmda-slope-policy-name*
      - **no slope-policy**
      - **stat-mode** {**v4-v6**}
      - **no stat-mode**
      - **wrr-weight** *weight*
      - **no wrr-weight**
- **ingress-qos**
  - **qos** *policy-id*
  - **no qos**

```
                        — queue queue-id
                            — stat-mode {v4-v6}
                            — no stat-mode
                        — policer policer-id
                            — stat-mode {v4-v6}
                            — no stat-mode
                — igmp-policy policy-name
                — no igmp-policy
                — ingress
                    — policer-control-policy policer-control-policy-name
                    — no policer-control-policy
                        — max-rate {rate | max}
                        — no max-rate
                        — priority-mbs-thresholds
                            — min-thresh-separation size [bytes | kilobytes]
                            — min-thresh-separation default
                            — no min-thresh-separation
                            — priority level
                                — mbs-contribution size [bytes | kilobytes]
                    — scheduler-policy scheduler-policy-name
                    — no scheduler-policy
                        — scheduler scheduler-name rate [cir cir-rate]
                        — no scheduler scheduler-name
                — pim-policy policy-name
                — no pim-policy policy-name
                — radius-accounting
                    — policy name1 [name2] [name3] [name4] [name5]
                    — no policy
                    — [no] session-optimized-stop
                — sla-profile-map
                    — entry key sla-profile-string sla-profile sla-profile-name
                    — no entry key sla-profile-string
                    — [no] use-direct-map-as-default
                — [no] secondary-shaper-hashing
                — sub-mcac-policy policy-name
                — no sub-mcac-policy
                — [no] vport-hashing
                — volume-stats-type {ip | default}
                — no volume-stats-type
```

9.25.1.24 IPoE Session Policy Commands

```
config
    — subscriber-mgmt
        — ipoe-session-policy policy-name [create]
        — no ipoe-session-policy policy-name
            — [no] circuit-id-from-auth
            — description description-string
            — no description
            — session-key sap mac [cid] [rid]
            — no session-key
```

```
— session-timeout timeout
— no session-timeout
```

9.25.1.25 SHCV Policy Commands

```
config
    — subscriber-mgmt
        — shcv-policy name [create]
        — no shcv-policy name
            — description description-string
            — no description
            — layer-3
                — source-ip-origin {interface | vrrp}
                — no source-ip-origin
            — periodic
                — action {alarm | remove}
                — no action
                — interval minutes
                — no interval
                — retry-count [retry-count]
                — [no] shutdown
                — timeout seconds
                — no timeout
            — trigger trigger-type
                — retry-count [retry-count]
                — no retry-count
                — [no] shutdown
                — timeout seconds
                — no timeout
            — vpls
                — source-ip ip-address
                — no source-ip
                — source-mac ieee-address
                — no source-mac
    — service
        — ies service-id
            — interface ip-int-name
                — shcv-policy-ipv4 name
                — no shcv-policy-ipv4
            — subscriber-interface ip-int-name
                — group-interface ip-int-name
                    — shcv-policy name
                    — no shcv-policy
                    — shcv-policy-ipv4 name
                    — no shcv-policy-ipv4
                    — shcv-policy-ipv6 name
                    — no shcv-policy-ipv6
        — vpls service-id
                — shcv-policy-ipv4 name
                — no shcv-policy-ipv4
            — sap
```

```
            — shcv-policy-ipv4 name
            — no shcv-policy-ipv4
      — vprn service-id
            — interface ip-int-name
                  — shcv-policy-ipv4 name
                  — no shcv-policy-ipv4
            — subscriber-interface ip-int-name
                  — group-interface ip-int-name
                        — shcv-policy name
                        — no shcv-policy
                        — shcv-policy-ipv4 name
                        — no shcv-policy-ipv4
                        — shcv-policy-ipv6 name
                        — no shcv-policy-ipv6
```

9.25.1.26 Subscriber Management Service Commands

9.25.1.26.1 VPLS Subscriber Management Configuration Commands

```
config
    — service
        — [no] vpls service-id [customer customer-id]  [vpn vpn-id] [m-vpls]
            — sap sap-id [split-horizon-group group-name] [create]
            — no sap sap-id
                — arp-host
                    — host-limit max-num-hosts
                    — no host-limit
                    — min-auth-interval min-auth-interval
                    — no min-auth-interval
                    — [no] shutdown
                — diameter-auth-policy name
                — no diameter-auth-policy
                — host-connectivity-verify source-ip ip-address [source-mac ieee-
                  address] [interval interval] [action {remove | alarm}] [timeout retry-
                  timeout] [retry-count count]
                — [no] sub-sla-mgmt
                    — def-app-profile default-app-profile-name
                    — no def-app-profile
                    — def-sla-profile default-sla-profile-name
                    — no def-sla-profile
                    — def-sub-id string sub-ident-string
                    — def-sub-id use-sap-id
                    — no def-sub-id
                    — def-sub-profile sub-profile-name
                    — no def-sub-profile
                    — [no] mac-da-hashing
                    — multi-sub-sap number-of-sub
                    — no multi-sub-sap
                    — [no] shutdown
```

```
                        — single-sub-parameters
                            — non-sub-traffic sub-profile sub-profile-name sla-
                                profile sla-profile-name [subscriber sub-ident-string]
                            — no non-sub-traffic
                            — [no] profiled-traffic-only
                        — sub-ident-policy sub-ident-policy-name
                        — no sub-ident-policy
                — sap sap-id [create] [capture-sap]
                    — ipoe-session
                        — description description-string
                        — no description
                        — ipoe-session-policy policy-name
                        — no ipoe-session-policy
                        — user-db local-user-db-name
                        — no user-db
                        — [no] shutdown
```

Capture SAP Configuration Commands

```
config
    — service
        — [no] vpls service-id [customer customer-id] [create]
            — sap sap-id [create] [capture-sap]
                — [no] allow-dot1q-msaps
                — authentication-policy name
                — no authentication-policy
                — cpu-protection policy-id [mac-monitoring] | [eth-cfm-monitoring
                    [aggregate] [car]]
                — no cpu-protection
                — [no] description
                — dhcp-python-policy policy-name
                — no dhcp-python-policy
                — dhcp-user-db local-user-db-name
                — no dhcp-user-db
                — dhcp6-user-db policy-name
                — no dhcp6-user-db
                — dhcp6-python-policy local-user-db-name
                — no dhcp6-python-policy
                — dist-cpu-protection policy-name
                — no dist-cpu-protection
                — dynamic-services
                    — dynamic-services-policy dynsvc-policy-name
                    — no dynamic-services-policy
                    — [no] shutdown
                — host-lockout-policy policy-name
                — no host-lockout-policy
                — msap-defaults
                    — group-interface ip-int-name
                    — no group-interface
                    — policy msap-policy-name
                    — no policy
```

```
                        — service service-id
                        — no service
                    — pppoe-policy pppoe-policy-name
                    — no pppoe-policy pppoe-policy-name
                    — pppoe-python-policy policy-name
                    — no pppoe-python-policy
                    — pppoe-user-db local-user-db-name
                    — no pppoe-user-db
                    — rtr-solicit-user-db local-user-db
                    — [no] rtr-solicit-user-db
                    — [no] shutdown
                    — track-srrp srrp-id
                    — no track-srrp
                    — trigger-packet [dhcp] [pppoe] [arp] [dhcp6] [ppp] [rtr-solicit] [data]
                    — no trigger-packet
```

9.25.1.26.2 Managed SAP Policy Commands

```
config
    — subscriber-mgmt
        — msap-policy msap-policy-name [create]
        — no msap-policy msap-policy-name
            — cpu-protection policy-id [mac-monitoring]
            — no cpu-protection
            — description description-string
            — no description
            — no dist-cpu-protection policy-name
            — ies-vprn-only-sap-parameters
                — anti-spoof {ip-mac | nh-mac}
                — no anti-spoof
                — egress
                    — qos policy-id
                    — no qos
                — ingress
                    — qos policy-id [shared-queuing | service-queuing]
                    — no qos
            — igmp-host-tracking
                — expiry-time expiry-time
                — no expiry-time
                — import policy-name
                — no import policy-name
                — max-num-groups max-num-groups
                — no max-num-groups
                — max-num-sources max-num-sources
                — no max-num-sources
                — max-num-grp-sources [max-num-grp-sources]
                — no max-num-grp-sources [max-num-grp-sources]
            — lag-link-map-profile link-map-profile-id
            — no lag-link-map-profile
            — [no] sticky-msaps
            — sub-sla-mgmt
```

```
        — def-app-profile app-profile-name
        — no def-app-profile
        — def-inter-dest-id {string string | use-top-q}
        — no def-inter-dest-id
        — def-sla-profile sla-profile-name
        — no def-sla-profile
        — def-sub-id use-auto-id
        — def-sub-id use-sap-id
        — def-sub-id string sub-id
        — no def-sub-id
        — def-sub-profile sub-profile-name
        — no def-sub-profile
        — multi-sub-sap [limit limit]
        — no multi-sub-sap
        — single-sub-parameters
            — non-sub-traffic sub-profile sub-profile-name sla-profile sla-
                profile-name [subscriber sub-ident-string] [app-profile
                app-profile-name]
            — no non-sub-traffic
            — [no] profiled-traffic-only
        — sub-ident-policy policy-name
        — no sub-ident-policy
    — vpls-only-sap-parameters
        — arp-host
            — host-limit max-num-hosts
            — no host-limit
            — min-auth-interval min-auth-interval
            — no min-auth-interval
        — arp-reply-agent [sub-ident]
        — no arp-reply-agent
        — egress
            — qos policy-id
            — no qos
        — ingress
            — qos policy-id [multipoint-shared | service-queuing]
            — no qos
        — igmp-snooping
            — [no] fast-leave
            — import policy-name
            — no import
            — last-member-query-interval interval
            — no last-member-query-interval
            — max-num-groups max-num-groups
            — no max-num-groups
            — mcac
                — if-policy mcac-if-policy-name
                — no if-policy
                — policy policy-name
                — no policy
                — mc-constraints
                    — level level-id bw bandwidth
                    — no level level-id
                    — number-down number-lag-port-down level level-
                        id
```

```
                        — no number-down number-lag-port-down
                        — use-lag-port-weight number-lag-port-down level
                             level-id
                        — no use-lag-port-weight
                    — unconstrained-bw bandwidth mandatory-bw
                         mandatory-bw
                    — no unconstrained-bw
                — mvr
                    — from-vpls service-id
                    — no from-vpls
                — query-interval seconds
                — no query-interval
                — query-response-interval seconds
                — no query-response-interval
                — robust-count robust-count
                — no robust-count
                — [no] send-queries
                — version version
                — no version
            — [no] mac-da-hashing
            — split-horizon-group group-name
            — no split-horizon-group
```

9.25.1.26.3 VPRN Subscriber Interface Commands

Note: The commands listed in this section apply only to the 7750 SR.

```
config
    — service
        — vprn service-id [customer customer-id]
        — no vprn service-id
            — subscriber-interface ip-int-name [create] fwd-service service-id fwd-
                 subscriber-interface fwd-int-name [wan-mode mode]
            — no subscriber-interface ip-int-name
                — address {ip-address/mask | ip-address netmask} [gw-ip-address ip-
                     address] [populate-host-routes]
                — no address
                — [no] allow-unmatching-subnets
                — [no] arp-host
                    — host-limit max-num-hosts
                    — no host-limit
                    — [no] shutdown
                — default-dns ip-address [secondary secondary-ip-address]
                — no default-dns
                — description description-string
                — no description
                — [no] export-host-routes
                — [no] ipoe-linking
```

```
            — [no] gratuitous-rtr-adv
        — [no] ipoe-session
            — session-limit session-limit
        — [no] ipv6
            — address ipv6-address/prefix-length [pd] [wan-host] [track-srrp
                  srrp-instance] [holdup-time milli-seconds]
            — no address ipv6-address/prefix-length
            — [no] allow-unmatching-subnets
            — default-dns ipv6-address [secondary ipv6-address]
            — no default-dns
            — delegated-prefix-length bits
            — delegated-prefix-length variable
            — no delegated-prefix-length
            — link-local-address ipv6-address
            — no link-local-address
            — subscriber-prefixes
                — prefix ipv6-address/prefix-length [pd] [wan-host] track-
                    srrp srrp-instance [holdup-time milli-seconds]
                — no prefix ipv6-address/prefix-length
                — prefix ipv6-address/prefix-length [pd] [wan-host]
            — [no] allow-multiple-wan-addresses
        — local-address-assignment
            — client-application [ppp-v4] [ipoe-v4]
            — no client-application
            — default-pool pool-name [secondary pool-name]
            — no default-pool
            — server server-name
            — no server
            — [no] shutdown
        — [no] private-retail-subnets
        — [no] shutdown
        — unnumbered [ip-int-name | ip-address]
        — no unnumbered
```

9.25.1.26.4 VPRN Group Interface Commands

Note: The commands listed in this section apply only to the 7750 SR.

```
config
    — service
        — vprn service-id [customer customer-id]
        — no vprn service-id
            — subscriber-interface ip-int-name [fwd-service service-id fwd-subscriber-
                interface ip-int-name] [create] [wan-mode mode]
            — no subscriber-interface ip-int-name
                — group-interface ip-int-name [create] [type]
                — no group-interface ip-int-name
                    — arp-host
```

- **host-limit** *max-num-hosts*
- **no host-limit**
- **min-auth-interval** *min-auth-interval*
- **no min-auth-interval**
- **sap-host-limit** *max-num-hosts-sap*
- **no sap-host-limit**
- [**no**] **shutdown**

— [**no**] **arp-populate**
— **arp-timeout** *seconds*
— **no arp-timeout**
— **authentication-policy** *name*
— **no authentication-policy**
— [**no**] **auto-reply**
- [**no**] **neighbor-solicitation**
- [**no**] **router-solicitation**

— **bonding-parameters**
- **connection** *connection-id* [**create**]
- **no connection***connection-id*
  - **service** *service-id*
  - **no service**
- **fpe** *fpe-id*
- **no fpe**
- **multicast**
  - **connection** *connection-id*
  - **connection use-incoming**
- [**no**] **shutdown**

— **cflowd-parameters**
- **sampling** {**unicast** | **multicast**} **type** {**acl** | **interface**} [**direction** {**ingress-only** | **egress-only** | **both**}]
- **no sampling** {**unicast** | **multicast**}

— **data-trigger**
- [**no**] **shutdown**

— **diameter-application-policy** *policy-name*
— **no diameter-application-policy**
— **diameter-auth-policy** *name*
— **no diameter-auth-policy**
— **description** *description-string*
— **no description**
— **diameter-application-policy** *policy-name*
— **no diameter-application-policy**
— [**no**] **enable-ingress-stats**
— **gtp-parameters**
- **fpe** *fpe-id*
- **no fpe**
- [**no**] **shutdown**

— **host-connectivity-verify** [**interval** *interval*] [**action** {**remove** | **alarm**}] [**timeout** *retry-timeout*] [**retry-count** *count*] [**family** *family*]
— **icmp**
- [**no**] **mask-reply**
- **param-problem** [*number seconds*]
- **no param-problem**
- **redirects** [*number seconds*]
- **no redirects**

- **ttl-expired** *number seconds*]
- **no ttl-expired**
- **unreachables** [*number seconds*]
- **no unreachables**

— **ingress**
- **policy-accounting** *template-name*
- **no policy-accounting**

— **ip-mtu** *octets*
— **no ip-mtu**
— **ipoe-linking**
- [**no**] **gratuitous-rtr-adv**
- [**no**] **shared-circuit-id**
- [**no**] **shutdown**

— **ipoe-session**
- **description** *description-string*
- **no description**
- **force-auth** [**cid-change**] [**rid-change**]
- **force-auth disabled**
- **no force-auth**
- **ipoe-session-policy** *policy-name*
- **no ipoe-session-policy**
- **min-auth-interval** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
- **min-auth-interval infinite**
- **no min-auth-interval**
- **radius-session-timeout** {**backwards-compatible** | **ignore** | **absolute**}
- **no radius-session-timeout**
- **sap-session-limit** *sap-session-limit*
- **no sap-session-limit**
- **session-limit** *session-limit*
- **no session-limit**
- **user-db** *local-user-db-name*
- **no user-db**
- [**no**] **shutdown**

— [**no**] **ipv6**
- **allow-multiple-wan-addresses**
- **no allow-multiple-wan-addresses**
- [**no**] **router-advertisements**
  - [**no**] **dns-options**
    - [**no**] **include-dns**
    - **rdnss-lifetime** *seconds*
    - **rdnss-lifetime infinite**
    - **no rdnss-lifetime**
- [**no**] **ipoe-bridged-mode**
- **router-solicit**
  - **inactivity-timer** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
  - **inactivity-timer infinite**
  - **no inactivity-timer**
  - **min-auth-interval** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
  - **no min-auth-interval**
  - [**no**] **shutdown**

```
            — [no] user-db
        — local-address-assignment
            — client-application [ppp-v4] [ipoe-v4]
            — no client-application
            — default-pool pool-name [secondary pool-name]
            — no default-pool
            — ipv6
                — client-application [ppp-slaac] [ipoe-wan] [ipoe-
                    slaac]
                — no client-application
                — server server-name
                — no server
            — server server-name
            — no server
            — [no] shutdown
        — [no] local-proxy-arp
        — [no] mac ieee-address
        — min-auth-interval min-auth-interval
        — no min-auth-interval
        — [no] oper-up-while-empty
        — policy-control diameter-policy-name
        — no policy-control
        — [no] proxy-arp-policy policy-name [policy-name]
        — [no] qos-route-lookup
        — redundant-interface red-ip-int-name
        — no redundant-interface
        — [no] remote-proxy-arp
        — [no] sap sap-id [create]
            — accounting-policy acct-policy-id
            — no accounting-policy [acct-policy-id]
            — anti-spoof {ip | ip-mac | nh-mac}
            — no anti-spoof
            — app-profile app-profile-name
            — no app-profile
            — atm
                — egress
                    — traffic-desc traffic-desc-profile-id
                    — no traffic-desc
                — encapsulation atm-encap-type
                — ingress
                    — traffic-desc traffic-desc-profile-id
                — oam
                    — [no] alarm-cells
                    — [no] periodic-loopback
            — [no] calling-station-id
            — [no] collect-stats
            — default-host ipv4-prefix/mask | ipv6-prefix/prefix-length
                next-hop ipv4-address | ipv6-address
            — no default-host ipv4-prefix/mask | ipv6-prefix/prefix-
                length
            — cpu-protection policy-id [mac-monitoring]
            — no cpu-protection
            — description description-string
            — no description
```

```
— egress
    — [no] agg-rate
        — [no] limit-unused-bandwidth
        — [no] queue-frame-based-accounting
        — rate rate
        — no rate
    — filter ip ip-filter-id
    — filter ipv6 ipv6-filter-id
    — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
    — no filter
    — [no] qinq-mark-top-only
    — [no] qos policy-id
    — scheduler-policy scheduler-policy-name
    — no scheduler-policy
— igmp-host-tracking
    — [no] disable-router-alert-check
    — expiry-time expiry-time
    — no expiry-time
    — import policy-name
    — no import
    — max-num-groups max-num-groups
    — no max-num-groups
    — max-num-sources max-num-sources
    — no max-num-sources
    — max-num-grp-sources [1 to 32000]
    — no max-num-grp-sources
— ingress
    — filter ip ip-filter-id
    — filter ipv6 ipv6-filter-id
    — no filter [ip ip-filter-id] [ipv6 ipv6-filter-id]
    — no filter
    — match-qinq-dot1p {top | bottom}
    — no match-qinq-dot1p
    — qos policy-id [shared-queuing]
    — no qos policy-id
    — scheduler-policy scheduler-policy-name
    — no scheduler-policy
— multi-service-site customer-site-name
— no multi-service-site
— static-host ip ip-prefix[/prefix-length] mac ieee-address]
    [create]
— no static-host ip-prefix[/prefix-length] mac ieee-
    address]
— no static-host all [force]
— no static-host ip ip-prefix[/prefix-length]
    — ancp-string ancp-string
    — no ancp-string
    — app-profile app-profile-name
    — no app-profile
    — inter-dest-id intermediate-destination-id
    — no inter-dest-id
    — managed-routes
        — route {ip-prefix/length | ip-prefix netmask}
            [create]
```

— **no route** {*ip-prefix*/*length* | *ip-prefix netmask*}
— **retail-svc-id** *service-id*
— **no retail-svc-id**
— **rip-policy** *policy-name*
— **no rip-policy**
— [**no**] **shutdown**
— **sla-profile** *sla-profile-name*
— **no sla-profile**
— **sub-profile** *sub-profile-name*
— **no sub-profile**
— **subscriber** *sub-ident*
— **no subscriber**
— [**no**] **subscriber-sap-id**
— **static-host-mgmt**
— [**no**] **mac-learning-options**
— [**no**] **data-triggered**
— [**no**] **single-mac**
— [**no**] **shutdown**
— [**no**] **sub-sla-mgmt**
— **def-app-profile** *default-app-profile-name*
— **no def-app-profile**
— **def-sla-profile** *default-sla-profile-name*
— **no def-sla-profile**
— **def-sub-id use-auto-id**
— **def-sub-id use-sap-id**
— **def-sub-id string** *sub-id*
— **no def-sub-id**
— **def-sub-profile** *sub-profile-name*
— **no def-sub-profile**
— **multi-sub-sap** *subscriber-limit*
— **no multi-sub-sap**
— [**no**] **shutdown**
— **single-sub-parameters**
— **non-sub-traffic sub-profile** *sub-profile-name* **sla-profile** *sla-profile-name* [**subscriber** *sub-ident-string*] [**app-profile** *app-profile-name*]
— **no non-sub-traffic**
— [**no**] **profiled-traffic-only**
— **sub-ident-policy** *sub-ident-policy-name*
— **no sub-ident-policy**
— [**no**] **shutdown**
— **srrp** *srrp-id* [**create**]
— **no srrp** *srrp-id*
— [**no**] **bfd-enable** *service-id* **interface** *interface-name* **dst-ip** *ip-address*
— **description** *description-string*
— **no description**
— **gw-mac** *mac-address*
— **no gw-mac**
— **keep-alive-interval** *interval*
— **no keep-alive-interval**
— **message-path** *sap-id*

- **monitor-oper-group** *group-name* **priority-step** *step*
- **no monitor-oper-group**
- [**no**] **one-garp-per-sap**
- **no message-path**
- [**no**] **policy** *vrrp-policy-id*
- [**no**] **preempt**
- **priority** *priority*
- **no priority**
- **send-fib-population-packets** *mode*
- **no send-fib-population-packets**
- [**no**] **shutdown**

— **wpp**

- [**no**] **enable-triggered-hosts**
- **initial-app-profile** *app-profile-name*
- **no initial-app-profile**
- **initial-sla-profile** *sla-profile-name*
- **no initial-sla-profile**
- **initial-sub-profile** *sub-profile-name*
- **no initial-sub-profile**
- **lease-time** [*lease-time*]
- **no lease-time**
- **portal router** *router-instance* **name** *wpp-portal-name*
- **no portal**
- **portal-group** *portal-group-name*
- **no portal-group**
- [**no**] **restore-disconnected**
- **user-db** *local-user-db-name*
- **no user-db**

9.25.1.26.5 IES Subscriber Management Configuration Commands

Note: The commands listed in this section apply only to the 7750 SR.

- IES Interface Commands
- IES Subscriber Interface Commands
- IES Group Interface Commands

IES Interface Commands

config

— **service**

— **ies** *service-id* [**customer** *customer-id*] [**vpn** *vpn-id*]

— [**no**] **interface** *ip-int-name* [**create**]

- **address** {*ip-address/mask* | *ip-address netmask*} [**broadcast all-ones** | **host-ones**] [**track-srrp** *srrp-instance*]
- **no address** [*ip-address/mask* | *ip-address netmask*]
- [**no**] **allow-directed-broadcasts**
- [**no**] **arp-populate**
- **arp-timeout** *seconds*
- **no arp-timeout**
- **cflowd-parameters**
  - **sampling** {**unicast** | **multicast**} **type** {**acl** | **interface**} [**direction** {**ingress-only** | **egress-only** | **both**}]
  - **sampling** {**unicast** | **multicast**}
- **description** *description-string*
- **no description**
- **host-connectivity-verify** [**source** {**vrrp** | **interface**}] [**interval** *interval*] [**action** {**remove** | **alarm**}] [**timeout** *retry-timeout*] [**retry-count** *count*]
- **icmp**
  - [**no**] **mask-reply**
  - **param-problem** [**number** *seconds*]
  - **no param-problem**
  - **redirects** [*number seconds*]
  - **no redirects**
  - **ttl-expired** *number seconds*]
  - **no ttl-expired**
  - **unreachables** [*number seconds*]
  - **no unreachables**
- **ipv6**
  - **address** *ipv6-address/prefix-length* [**pd**] [**wan-host**] [**track-srrp** *srrp-instance*] [**holdup-time** *milli-seconds*]
  - **no address** *ipv6-address/prefix-length*
  - **bfd** *transmit-interval* [**receive** *receive-interval*] [**multiplier** *multiplier* [**echo-receive** *echo-interval*] [**type** *cpm-np*]
  - **no bfd**
  - [**no**] **dad-disable**
  - [**no**] **dhcp6-relay**
  - [**no**] **dhcp6-server**
  - **icmp6**
    - **packet-too-big** [*number seconds*]
    - **no packet-too-big**
    - **param-problem** [*number seconds*]
    - **no param-problem**
    - **redirects** [*number seconds*]
    - **no redirects**
    - **time-exceeded** [*number seconds*]
    - **no time-exceeded**
    - **unreachables** [*number seconds*]
    - **no unreachables**
  - **link-local-address** *ipv6-address*
  - **no link-local-address**
  - **local-dhcp-server** *local-server-name*
  - **no local-dhcp-server**
  - [**no**] **local-proxy-nd**
  - **neighbor** *ipv6-address mac-address*
  - **no neighbor** *ipv6-address*

- **neighbor-limit** *limit* [**log-only**] [**threshold** *percent*]
- **no neighbor-limit**
- **proxy-nd-policy** *policy-name* [*policy-name*]
- **no proxy-nd-policy**
- **qos-route-lookup** [**source** | **destination**]
- **no qos-route-lookup**
- **reachable-time** *seconds*
- **no reachable-time**
- [**no**] **secure-nd**
- **stale-time** *seconds*
- **no stale-time**
- **tcp-mss** *mss-value*
- **tcp-mss**
- [**no**] **urpf-check**
  - [**no**] **ignore-default**
  - **mode** {**strict** | **loose** | **strict-no-ecmp**}

- [**no**] **local-proxy-arp**
- [**no**] **loopback**
- [**no**] **mac** *ieee-address*
- [**no**] **proxy-arp-policy** *policy-name* [*policy-name*]
- [**no**] **proxy-arp-policy**
- [**no**] **sap** *sap-id* **— Refer to the 7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN for information about SAP CLI commands and descriptions.**

9.25.1.26.6 IES Subscriber Interface Commands

Note: The commands listed in this section apply only to the 7750 SR.

```
config
    — service
        — ies service-id [customer customer-id] [vpn vpn-id]
            — subscriber-interface ip-int-name [create] fwd-service service-id fwd-
                subscriber-interface fwd-int-name [wan-mode mode]
            — no subscriber-interface ip-int-name
                — address ip-address/mask [netmask] [gw-ip-address gw-ip-address]
                    [populate-host-routes] [track-srrp srrp-instance] [holdup-time
                    milli-seconds]
                — address ip-address/mask [netmask]
                — [no] allow-unmatching-subnets
                — description description-string
                — no description
                — [no] export-host-routes
                — ipoe-linking
                    — [no] gratuitous-rtr-adv
                    — [no] shared-circuit-id
                    — [no] shutdown
                — [no] ipoe-session
                    — session-limit session-limit
                    — no session-limit
                — ipv6
                    — address ip-address/mask [netmask] [gw-ip-address gw-ip-
                        address] [populate-host-routes] [track-srrp srrp-
                        instance] [holdup-time milli-seconds]
                    — no address ip-address/mask [netmask]
                    — [no] allow-multiple-wan-addresses
                    — [no] allow-unmatching-prefixes
                    — default-dns ipv6-address [secondary ipv6-address]
                    — no default-dns
                    — delegated-prefix-length bits
                    — delegated-prefix-length variable
                    — no delegated-prefix-length
                    — [no] ipoe-bridged-mode
                    — link-local-address ipv6-address
                    — no link-local-address
                    — [no] router-advertisements
                        — current-hop-limit limit
                        — no current-hop-limit
                        — [no] dns-options
                            — [no] include-dns
                            — rdnss-lifetime seconds
                            — rdnss-lifetime infinite
                            — no rdnss-lifetime
                        — force-mcast [ip] [mac]
                        — no force-mcast
```

```
                    — [no] managed-configuration
                    — max-advertisement seconds
                    — no max-advertisement
                    — min-advertisement seconds
                    — no min-advertisement
                    — mtu bytes
                    — no mtu
                    — [no] other-stateful-configuration
                    — [no] prefix-options
                            — [no] autonomous
                        — [no] on-link
                        — preferred-lifetime seconds
                        — preferred-lifetime infinite
                        — no preferred-lifetime
                        — valid-lifetime seconds
                        — valid-lifetime infinite
                        — no valid-lifetime
                        — reachable-time milli-seconds
                        — no reachable-time
                        — retransmit-time milli-seconds
                        — no retransmit-time
                        — router-lifetime seconds
                        — router-lifetime no-default-router
                        — no router-lifetime
                        — [no] shutdown
                — [no] router-solicit
                    — inactivity-timer [days days] [hrs hours] [min minutes]
                      [sec seconds]
                    — inactivity-timer infinite
                    — no inactivity-timer
                — subscriber-prefixes
                    — prefix ipv6-address/prefix-length [pd] [wan-host] track-
                      srrp srrp-instance [holdup-time milli-seconds]
                    — no prefix ipv6-address/prefix-length
            — local-address-assignment
                — client-application [ppp-v4] [ipoe-v4]
                — no client-application
                — default-pool pool-name [secondary pool-name]
                — no default-pool
                — server server-name
                — no server
                — [no] shutdown
                — ipv6
                    — client-application [ppp-slaac] [ipoe-wan] [ipoe-slaac]
                    — no client-application
                    — server server-name
                    — no server
            — unnumbered [ip-int-name | ip-address]
            — no unnumbered
```

9.25.1.26.7 IES Group Interface Commands

Note: The commands listed in this section apply only to the 7750 SR.

```
config
    — service
        — ies service-id [customer customer-id] [vpn vpn-id]
            — [no] subscriber-interface ip-int-name
                — [no] group-interface ip-int-name
                    — arp-host
                        — host-limit max-num-hosts
                        — no host-limit
                        — min-auth-interval min-auth-interval
                        — no min-auth-interval
                        — sap-host-limit max-num-hosts-sap
                        — no sap-host-limit
                        — [no] shutdown
                    — [no] arp-populate
                    — arp-timeout seconds
                    — no arp-timeout
                    — authentication-policy name
                    — no authentication-policy
                    — bonding-parameters
                        — connection connection-id [create]
                        — no connection
                            — service service-id
                            — no service
                        — fpe fpe-id
                        — no fpe
                        — multicast
                            — connection connection-id
                            — connection use-incoming
                        — [no] shutdown
                    — cflowd-parameters
                        — sampling {unicast | multicast} type {acl | interface}
                          [direction {ingress-only | egress-only | both}]
                        — no sampling {unicast | multicast}
                    — data-trigger
                        — [no] shutdown
                    — description description-string
                    — no description
                    — diameter-application-policy policy-name
                    — no diameter-application-policy
                    — diameter-auth-policy name
                    — no diameter-auth-policy
                    — gtp-parameters
                        — fpe fpe-id
                        — no fpe
                        — [no] shutdown
```

- **host-connectivity-verify** [**interval** *interval*] [**action** {**remove** | **alarm**}] [**timeout** *retry-timeout*] [**retry-count** *count*] [**family** *family*]
- **icmp**
  - [**no**] **mask-reply**
  - **param-problem** [*number seconds*]
  - **no param-problem**
  - **redirects** [*number seconds*]
  - **no redirects**
  - **ttl-expired** *number seconds*]
  - **no ttl-expired**
  - **unreachables** [*number seconds*]
  - **no unreachables**
- **ip-mtu** *octets*
- **no ip-mtu**
- **ipoe-linking**
  - [**no**] **gratuitous-rtr-adv**
  - [**no**] **shared-circuit-id**
  - [**no**] **shutdown**
- **ipoe-session**
  - **description** *description-string*
  - **no description**
  - **force-auth** [**cid-change**] [**rid-change**]
  - **force-auth disabled**
  - **no force-auth**
  - **ipoe-session-policy** *policy-name*
  - **no ipoe-session-policy**
  - **min-auth-interval** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
  - **min-auth-interval infinite**
  - **no min-auth-interval**
  - **radius-session-timeout** {**backwards-compatible** | **ignore** | **absolute**}
  - **no radius-session-timeout**
  - **sap-session-limit** *sap-session-limit*
  - **no sap-session-limit**
  - **session-limit** *session-limit*
  - **no session-limit**
  - [**no**] **shutdown**
  - **user-db** *local-user-db-name*
  - **no user-db**
- **ipv6**
  - [**no**] **allow-multiple-wan-addresses**
  - [**no**] **auto-reply**
    - [**no**] **neighbor-solicitation**
    - [**no**] **router-solicitation**
  - [**no**] **router-advertisements**
    - [**no**] **dns-options**
      - [**no**] **include-dns**
      - **rdnss-lifetime** *seconds*
      - **rdnss-lifetime infinite**
      - **no rdnss-lifetime**
- **local-address-assignment**
  - **client-application** [**ppp-v4**] [**ipoe-v4**]

```
            — no client-application
            — default-pool pool-name [secondary pool-name]
            — no default-pool
            — [no] ipv6
                — client-application [ppp-slaac] [ipoe-wan] [ipoe-
                    slaac]
                — no client-application
                — server server-name
                — no server
            — server server-name
            — no server
            — [no] shutdown
        — [no] local-proxy-arp
        — [no] mac ieee-address
        — [no] oper-up-while-empty
        — [no] proxy-arp-policy policy-name [policy-name]
        — [no] proxy-arp-policy
        — [no] sap sap-id
            — accounting-policy acct-policy-id
            — no accounting-policy [acct-policy-id]
            — anti-spoof {ip | ip-mac | nh-mac}
            — no anti-spoof
            — atm
                — egress
                    — traffic-desc traffic-desc-profile-id
                    — no traffic-desc
                — encapsulation atm-encap-type
                — ingress
                    — traffic-desc traffic-desc-profile-id
                — oam
                    — [no] alarm-cells
                    — [no] periodic-loopback
            — [no] calling-station-id
            — [no] collect-stats
            — cpu-protection policy-id [mac-monitoring] | [eth-cfm-
                monitoring [aggregate] [car]]
            — no cpu-protection
            — default-host ip-address/mask next-hop next-hop-ip
            — no default-host ip-address/mask
            — description description-string
            — no description
            — egress
                — [no] agg-rate
                    — [no] limit-unused-bandwidth
                    — [no] queue-frame-based-accounting
                    — rate kilobits-per-second
                    — no rate
                — filter {ip ip-filter-id | ipv6 ipv6-filter-id}
                — [no] qinq-mark-top-only
                — no filter [{ip ip-filter-id | ipv6 ipv6-filter-id}]
                — qos policy-id
                — no qos
                — policer-control-policy policy-name
                — no policer-control-policy
```

```
            — scheduler-policy scheduler-policy-name
            — no scheduler-policy
        — ingress
            — filter {ip ip-filter-id | ipv6 ipv6-filter-id}
            — no filter [{ip ip-filter-id | ipv6 ipv6-filter-id}]
            — match-qinq-dot1p {top | bottom}
            — no match-qinq-dot1p
            — qos policy-id [shared-queuing]
            — no qos
            — policer-control-policy policy-name
            — no policer-control-policy
            — scheduler-policy scheduler-policy-name
            — no scheduler-policy
        — static-host ip ip-address [mac ieee-address] [create]
        — static-host mac ieee-address [create]
        — no static-host [ip ip-address] mac ieee-address
        — no static-host all [force]
        — no static-host ip ip-address
            — ancp-string ancp-string
            — no ancp-string
            — app-profile app-profile-name
            — no app-profile
            — inter-dest-id intermediate-destination-id
            — no inter-dest-id
            — mac-linking ip-address
            — no mac-linking
            — managed-routes
                — [no] route {ip-prefix/length | ip-prefix
                    netmask}
                — [no] route ipv6-prefix/prefix-length [metric
                    metric-value]
            — rip-policy policy-name
            — no rip-policy
            — [no] shutdown
            — sla-profile sla-profile-name
            — no sla-profile
            — sub-profile sub-profile-name
            — no sub-profile
            — subscriber sub-ident
            — no subscriber
            — [no] subscriber-sap-id
            — [no] shutdown
            — static-host-mgmt
                — [no] mac-learning-options
                    — [no] data-triggered
                    — [no] single-mac
    — [no] sub-sla-mgmt
        — def-app-profile default-app-profile-name
        — no def-app-profile
        — def-sla-profile default-sla-profile-name
        — no def-sla-profile
        — def-sub-id use-auto-id
        — def-sub-id use-sap-id
        — def-sub-id string sub-id
```

```
            — no def-sub-id
            — def-sub-profile sub-profile-name
            — no def-sub-profile
            — multi-sub-sap subscriber-limit
            — no multi-sub-sap
            — [no] shutdown
            — single-sub-parameters
                — non-sub-traffic sub-profile sub-profile-name
                    sla-profile sla-profile-name [subscriber sub-
                    ident-string] [app-profile app-profile-name]
                — no non-sub-traffic
                — [no] profiled-traffic-only
            — sub-ident-policy sub-ident-policy-name
            — no sub-ident-policy
        — [no] shutdown
        — srrp srrp-id [create]
        — no srrp srrp-id
            — [no] bfd-enable service-id interface interface-name dst-
                ip ip-address
            — description description-string
            — no description
            — gw-mac mac-address
            — no gw-mac
            — keep-alive-interval interval
            — no keep-alive-interval
            — message-path sap-id
            — monitor-oper-group group-name priority-step step
            — no monitor-oper-group
            — [no] one-garp-per-sap
            — no message-path
            — [no] policy vrrp-policy-id
            — [no] preempt
            — priority priority
            — no priority
            — send-fib-population-packets mode
            — no send-fib-population-packets
            — [no] shutdown
        — srrp-enabled-routing [hold-time hold-time]
        — no srrp-enabled-routing
        — tos-marking-state {trusted | untrusted}
        — no tos-marking-state
        — [no] ignore-default
        — wpp
            — [no] enable-triggered-hosts
            — initial-app-profile app-profile-name
            — no initial-app-profile
            — initial-sla-profile sla-profile-name
            — no initial-sla-profile
            — initial-sub-profile sub-profile-name
            — no initial-sub-profile
            — lease-time [lease-time]
            — no lease-time
            — portal router router-instance name wpp-portal-name
            — no portal
```

- **portal-group** *portal-group-name*
- **no portal-group**
- [**no**] **restore-disconnected**
- **user-db** *local-user-db-name*
- **no user-db**

9.25.1.26.8 Service Subscriber Interface, Group Interface IPoE Command

VPRN Group Interface Commands

config
- **service**
  - **vprn** *service-id* [**customer** *customer-id*]
  - **no vprn** *service-id*
  - **ies** *service-id* [**customer** *customer-id*]
  - **no ies** *service-id*
    - **subscriber-interface** *ip-int-name* [**fwd-service** *service-id* **fwd-subscriber-interface** *ip-int-name*] [**create**] [**wan-mode** *mode*]
    - **no subscriber-interface** *ip-int-name*
      - **group-interface** *ip-int-name* [**create**]
      - **group-interface** *ip-int-name* [**create**] **lns**
      - **group-interface** *ip-int-name* [**create**] **wlangw**
      - **no group-interface** *ip-int-name*
        - [**no**] **ipoe-session**
          - **description** *description-string*
          - **no description**
          - **force-auth** [**cid-change**] [**rid-change**]
          - **force-auth disabled**
          - **no force-auth**
          - **ipoe-session-policy** *policy-name*
          - **no ipoe-session-policy**
          - **min-auth-interval** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec seconds**]
          - **min-auth-interval infinite**
          - **no min-auth-interval**
          - **sap-session-limit** *sap-session-limit*
          - **no sap-session-limit**
          - **session-limit** *session-limit*
          - **no session-limit**
          - [**no**] **shutdown**
          - **user-db** *local-user-db-name*
          - **no user-db**

9.25.1.26.9 RIP Commands

configure
- **subscriber-mgmt**

```
— rip-policy policy-name [create]
— no rip-policy policy-name
    — authentication-key authentication-key | hash-key [hash | hash2]
    — no authentication-key
    — authentication-type {none | password | message-digest | message-digest-
        20}
    — no authentication-type
    — description description-string
    — no description
— local-user-db local-user-db-name [create]
— no local-user-db local-user-db-name
    — ipoe
        — host host-name [create]
        — no host host-name
            — rip-policy policy-name
            — no rip-policy
    — ppp
        — host host-name [create]
        — no host host-name
            — rip-policy policy-name
            — no rip-policy

configure
    — service
        — {ies | vprn}
            — [no] subscriber-interface ip-int-name
                — [no] group-interface ip-int-name
                    — [no] sap sap-id
                        — static-host ip ip/did-address [mac ieee-address]
                            [create]
                        — static-host mac ieee-address [create]
                        — no static-host [ip ip-address] mac ieee-address
                        — no static-host all [force]
                        — no static-host ip ip-address
                            — rip-policy policy-name
                            — no rip-policy
        — vprn
            — [no] rip
                — [no] group name
                    — [no] neighbor ip-int-name

configure
    — router
        — [no] rip
            — [no] group name
                — [no] neighbor ip-int-name
```

9.25.1.26.10 VPort Commands

Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Interface Configuration Guide* for further information on card, Media Dependent Adapter (MDA), MCM (MDA Carrier Module), CMA (Compact Media Adapter) and port provisioning.

config
- **port** *port-id*
  - **ethernet**
    - **access**
      - **egress**
        - **vport** *vport-name* [**create**]
        - **no vport** *vport-name*
          - **agg-rate**
            - **rate** *kilobits-per-second*
          - **agg-rate-limit** *limit*
          - **no agg-rate-limit**
          - **description** *description-string*
          - **no description**
          - [**no**] **egress-rate-modify**
          - **host-match dest** *string* [**create**]
          - **no host-match dest** *string*
          - **port-scheduler-policy** *port-scheduling-policy-name*
          - **no port-scheduler-policy**
          - **scheduler-policy** *scheduler-policy-name*
          - **no scheduler-policy**

9.25.1.26.11 Redundant Interface Commands

Note: The commands listed in this section apply only to the 7750 SR.

config
- **service**
  - **ies**
    - [**no**] **redundant-interface** *ip-int-name*
      - **address** {*ip-address/mask* | *ip-address netmask*} [**remote-ip** *ip-address*]
      - **no address**
      - [**no**] **description** *description-string*
      - [**no**] **shutdown**
      - [**no**] **spoke-sdp** *sdp-id:vc-id* [**create**]
        - **egress**
          - **filter** [**ip** *ip-filter-id*]
          - **vc-label** *egress-vc-label*
          - **no vc-label** [*egress-vc-label*]
        - **ingress**
          - **filter** [**ip** *ip-filter-id*]

```
                                    — no filter
                                    — vc-label ingress-vc-label
                                    — no vc-label [ingress-vc-label]
                                — [no] shutdown

config
    — service
        — vprn
            — [no] redundant-interface ip-int-name
                — address {ip-address/mask | ip-address netmask} [remote-ip ip-
                    address]
                — no address
                — [no] description description-string
                — [no] shutdown
                — [no] spoke-sdp sdp-id:vc-id
                    — egress
                        — filter [ip ip-filter-id]
                        — vc-label egress-vc-label
                        — no vc-label [egress-vc-label]
                    — ingress
                        — filter [ip ip-filter-id]
                        — no filter
                        — vc-label ingress-vc-label
                        — no vc-label [ingress-vc-label]
                    — [no] shutdown

config
    — service
        — ies service-id [customer customer-id] [vpn vpn-id]
        — vprn service-id [customer customer-id]
        — no vprn service-id
            — interface ip-int-name [create] [tunnel]
            — no interface ip-int-name
                — [no] ipv6
                    — [no] ignore-default
                        — mode {strict | loose | strict-no-ecmp}
                        — no mode
            — [no] subscriber-interface ip-int-name
                — group-interface ip-int-name [create]
                — group-interface ip-int-name [create] lns
                — group-interface ip-int-name [create] wlangw
                — no group-interface ip-int-name
                    — [no] ipv6
                        — [no] ignore-default
                            — mode {strict | loose | strict-no-ecmp}
                            — no mode
```

9.25.1.26.12 Wireless Portal Protocol (WPP) Commands

```
config
    — aaa
```

```
            — wpp
                — portal-groups
                    — portal-group portal-group-name [create]
                    — no portal-group portal-group-name
                        — description description-string
                        — no description
                        — [no] portal router router-instance name wpp-portal-name
                        — [no] shutdown

config
    — router
        — wpp
            — portals
                — portal wpp-portal-name address ip-address [create]
                — portal wpp-portal-name
                — no portal wpp-portal-name
                    — ack-auth-retry-count [value]
                    — no ack-auth-retry-count
                    — ntf-logout-retry-count [value]
                    — no ntf-logout-retry-count
                    — port-format formatting
                    — no port-format
                    — retry-interval milliseconds
                    — no retry-interval
                    — secret secret [hash | hash2]
                    — no secret
                    — [no] shutdown
                    — version version
                    — no version
                — [no] shutdown

config
    — service
        — vprn service-id [customer customer-id]
        — no vprn service-id
            — wpp
                — portals
                    — portal wpp-portal-name address ip-address [create]
                    — portal wpp-portal-name
                    — no portal wpp-portal-name
                        — ack-auth-retry-count [value]
                        — no ack-auth-retry-count
                        — ntf-logout-retry-count [value]
                        — no ntf-logout-retry-count
                        — retry-interval milliseconds
                        — no retry-interval
                        — secret secret [hash | hash2]
                        — no secret
                        — [no] shutdown
                        — version version
                        — no version
                — [no] shutdown
```

```
config
    — service
        — ies service-id [customer customer-id] [vpn vpn-id]
        — vprn service-id [customer customer-id]
        — no vprn service-id
            — subscriber-interface ip-int-name [create]
            — no subscriber-interface ip-int-name
                — group-interface ip-int-name [create]
                — no group-interface ip-int-name
                    — wpp
                        — [no] enable-triggered-hosts
                        — initial-app-profile app-profile-name
                        — no initial-app-profile
                        — initial-sla-profile sla-profile-name
                        — no initial-sla-profile
                        — initial-sub-profile sub-profile-name
                        — no initial-sub-profile
                        — lease-time [lease-time]
                        — no lease-time
                        — portal router router-instance name wpp-portal-name
                        — no portal
                        — portal-group portal-group-name
                        — no portal-group
                        — [no] restore-disconnected
                        — [no] shutdown
                        — user-db local-user-db-name
                        — no user-db
                        — [no] shutdown

config
    — subscriber-mgmt
        — local-user-db local-user-db-name [create]
        — no local-user-db local-user-db-name
            — ipoe
                — host host-name [create]
                — no host host-name
                    — wpp
                        — initial-app-profile app-profile-name
                        — no initial-app-profile
                        — initial-sla-profile sla-profile-name
                        — no initial-sla-profile
                        — initial-sub-profile sub-profile-name
                        — no initial-sub-profile
                        — portal router router-instance name wpp-portal-name
                        — no portal
                        — portal-group portal-group-name
                        — no portal-group
                        — restore-disconnected {restore | no-restore}
                        — no restore-disconnected
```

9.25.1.26.13 Multiple PPPoE Session QoS Commands

```
config
    — qos
        — sap-egress policy-id
            — parent-location {default | sla}
            — no parent-location
    — scheduler-policy name [create]
        — tier 1
            — parent-location {none | sub | vport}
            — no parent-location

config
    — port
        — ethernet
            — access
                — egress
                    — vport name
                        — port-scheduler-policy port-scheduler-policy-name
                        — no port-scheduler-policy
                        — scheduler-policy scheduler-policy-name
                        — no scheduler-policy

config
    — subscr-mgmt
        — sla-prof
            — egress
                — scheduler-policy scheduler-policy-name
                    — scheduler-policy scheduler-name rate pir-rate [cir cir-
                      rate]

configure
    — subscriber-mgmt
        — authentication-policy name
            — include-radius-attribute
                — [no] sap-session-index
```

9.25.1.26.14 Multicast Listener Discovery (MLD) Commands

For more information about MLD commands, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Multicast Routing Protocols Guide*.

```
config
    — subscr-mgmt
        — mld-policy mld-policy-name [create]
        — no mld-policy mld-policy-name
            — description description-string
            — no description
            — [no] disable-router-alert-check
            — egress-rate-modify agg-rate-limit
```

- **egress-rate-modify scheduler** *scheduler-name*
- **no egress-rate-modify**
- [**no**] **fast-leave**
- **import** *policy-name*
- **no import**
- **max-num-groups** *max-num-groups*
- **no max-num-groups**
- **max-num-grp-sources** [*max-num-grp-sources*]
- **no max-num-grp-sources**
- **max-num-sources** *max-num-sources*
- **no max-num-sources**
- [**no**] **per-host-replication** [**uni-mac** | **mcast-mac**]
- **query-interval** *seconds*
- **no query-interval**
- **query-last-listener-interval** *seconds*
- **no query-last-listener-interval**
- **query-response-interval** *seconds*
- **no query-response-interval**
- **redirection-policy** *policy-name*
- **no redirection-policy**
- **static**
    - [**no**] **group** *grp-ipv6-address*
        - [**no**] **source** *ipv6-address*
        - [**no**] **starg**
- **version** *version*
- **no version**

9.25.1.27 Show Commands

Note:

- Refer to the 7450 ESS, 7750 SR, 7950 XRS, and VSR Quality of Service Guide for information about QoS commands, descriptions, CLI syntax, and usage.

show
- **aaa**
  - **diameter-node** *origin-host-string* **peer** *dest-host-string*
  - **diameter-node** *origin-host-string* **peer** *dest-host-string* **statistics**
  - **diameter-node** *origin-host-string* **peers**
  - **diameter-node** *origin-host-string* **routing-table**
  - **diameter-peer-policy** [*peer-policy-name*]
  - **diameter-peer-policy** *peer-policy-name* **peer** *peer-name* [**statistics**]
  - **diameter-peer-policy** *peer-policy-name* **client-side-peer-ip** *ip-address* **port** *tcp-port* [**statistics**]
  - **isa-radius-policy** *policy-name*
  - **isa-radius-policy** *policy-name* **associations**
  - **radius-configuration**
  - **radius-server-policy** *policy-name* [**acct-on-off**]
  - **radius-server-policy** *policy-name* **associations**
  - **radius-server-policy** *policy-name* **msg-buffer-stats**
  - **radius-server-policy** *policy-name* **statistics**
  - **radius-server-policy** [**acct-on-off**]
  - **wpp portal-group** *name*
  - **wpp** [**portal-group** *name*] [**host** *ip-address*] **hosts**

show
- **router**
  - **radius-proxy-server** *server-name*
  - **radius-proxy-server** *server-name* **cache**
  - **radius-proxy-server** *server-name* **cache hex-key** *hex-string*
  - **radius-proxy-server** *server-name* **cache string-key** *string*
  - **radius-proxy-server** *server-name* **cache summary**
  - **radius-proxy-server** *server-name* **statistics**
  - **radius-proxy-server**
  - **wpp**
  - **wpp** [**portal** *wpp-portal-name*] [**host** *ip-address*] **hosts**
  - **wpp portal** *wpp-portal-name*
  - **wpp statistics**

show
- **service**
  - **bonding** [**detail**]
  - **bonding bonding-id** *bonding-id* [**detail**]
  - **id** *service-id*
    - **gsmp**
    - **managed-hosts type** {**aaa** | **data-triggered**} [**ip** *ip-address*] [**mac** *ieee-address*]
    - **slaac**

— **host** [**detail**] [**wholesaler** *service-id*] [**session** {**none** | **ipoe** | **ppp**}] [**router-advertisement-policy** *policy-name*]
— **host interface** *interface-name* [**detail**] [**wholesaler** *service-id*] [**session** {**none** | **ipoe** | **ppp**}] [**router-advertisement-policy** *policy-name*]
— **host mac** *ieee-address* [**detail**] [**wholesaler** *service-id*] [**session** {**none host** | **ipoe** | **ppp**}] [**router-advertisement-policy** *policy-name*]
— **host ipv6-address** *ipv6-prefix* [**detail**] [**wholesaler** *service-id*] [**session** {**none** | **ipoe** | **ppp**}] [**router-advertisement-policy** *policy-name*]
— **host sap** *sap-id* [**detail**] [**wholesaler** *service-id*] [**session** {**none** | **ipoe** | **ppp**}] [**router-advertisement-policy** *policy-name*]
— **summary**
— **subscriber-hosts** [**sap** *sap-id*] [**ip** *ip-prefix/prefix-length*] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**app-profile** *app-profile-name*] [**wholesaler** *service-id*] [**address-origin** *address-origin*] [**detail**] [**statistics**] [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]
— **subscriber-using** [**service-id** *service-id*] [**sap-id** *sap-id*] [**interface** *ip-int-name*] [**ip** *ip-address*[/*mask*]] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**app-profile** *app-profile-name*] [**port** *port-id*] [**no-inter-dest-id** | **inter-dest-id** *intermediate-destination-id*]

show
— **subscriber-mgmt**
— **accu-stats** *subscriber subscriber-id*
— **accu-stats-policy** *policy-name* [**associations**]
— **accu-stats-subscribers** [**subscriber** *sub-ident*] [**sub-profile** *name*]
— **ancp-string** [*policy-name*]
— **ancp-string** *policy-name* **association**
— **ancp-string**
— **ancp-string** *ancp-string*
— **ancp-string customer** *customer-id* **site** *customer-site-name*
— **ancp-string sap** *sap-id*
— **authentication** *policy-name* **association**
— **authentication** [*policy-name*]
— **authentication** [*policy-name*] **statistics**
— **authentication coa-statistics**
— **category-map** [*category-map*]
— **category-map** *category-map* **subscribers**
— **diameter-application-policy** [name]
— **diameter-session**
— **ccrt-replay** [**session-id** *session-id*] [**diameter-application-policy** *name*]
— **ccrt-replay diameter-application-policy** *name* **summary**
— **ccrt-replay summary**
— **errors**
— **mac all** | *ieee-address* | **unknown**
— **sap all** | *sap-id*
— **sdp all** | *sdp-id:vc-id*
— **unknown-origin**
— **explicit-subscriber-map**
— **host-lockout-policy**
— **host-lockout-policy** *policy-name* **association**
— **host-lockout-policy** *policy-name*
— **host-lockout-policy** *policy-name lockout-state* **sap** *sap-id* [**summary**]

- **host-lockout-policy** *policy-name lockout-state* [**circuit-id** *circuit-id*] [**sap** *sap-id*] [**summary**]
- **host-lockout-policy** *policy-name lockout-state* [**mac** *ieee-address*] [**sap** *sap-id*] [**summary**]
- **host-lockout-policy** *policy-name lockout-state* [**remote-id** *remote-id*] [**sap** *sap-id*] [**summary**]
- **igmp-policy**
- **igmp-policy** *policy-name* **association**
- **igmp-policy** *policy-name*
- **ipoe-session-policy** *ipoe-session-policy-name* **association**
- **ipoe-session-policy**
- **local-user-db** *local-user-db-name* **association** [**dhcp**] [**ppp**] [**l2tp**] [**radius**] [**pppoe**] [**dhcp6**] [**capture-sap**] [**rtr-solicit**] [**wpp**] [**ipoe**]
- **local-user-db** *local-user-db-name* **ipoe-host** *ipoe-host-name*
- **local-user-db** *local-user-db-name* **ipoe-all-hosts**
- **local-user-db** *local-user-db-name* **ipoe-unmatched-hosts**
- **local-user-db** [*local-user-db-name*]
- **local-user-db** *local-user-db-name* **ppp-all-hosts**
- **local-user-db** *local-user-db-name* **ppp-host** *pppoe-host-name*
- **local-user-db** *local-user-db-name* **ppp-unmatched-hosts**
- **msap-policy** [*msap-policy-name* [**association**]]
- **pcc-rule**
- **pcc-rule monitoring-key** *key* **detail**
- **pcc-rule rule-id** *id* **detail**
- **pcc-rule rule-name** *rule-name*
- **pcc-rule rule-name** *rule-name* **detail**
- **pcc-rule summary**
- **pcc-rule monitoring-key** *key*
- **radius-accounting-policy** *name* **association**
- **radius-accounting-policy** [*name*]
- **radius-accounting-policy** *name* **statistics**
- **sla-profile** [*sla-profile-name* [**association**]]
- **statistics iom** (*slot* | **all**) [**host** | **session** | **subscriber** | **summary**] [**non-zero-value-only**]
- **statistics mda** (*mda* | **all**) [**host** | **session** | **subscriber** | **summary**] [**non-zero-value-only**]
- **statistics port** (*port-id* | **all**) [**host** | **session** | **subscriber** | **summary**] [**non-zero-value-only**]
- **statistics pw-port** (*pw-port* | **all**) [**host** | **session** | **subscriber** | **summary**] [**non-zero-value-only**]
- **statistics sla-profiles**
- **statistics sub-profiles**
- **statistics system** [**host** | **session** | **subscriber** | **summary**] [**non-zero-value-only**]
- **steering-profile** *steering-profile-name*
- **sub-ident-policy** [*sub-ident-policy-name* [**association**]]
- **sub-ident-policy** *sub-ident-policy-name* **script** {**primary** | **secondary** | **tertiary**}
- **sub-profile** [*sub-profile-name* [**association**]]

show

- **call-trace**
  - **files** [**running** | **finished**] [**detail**]
  - **ipoe**
    - **session** [**name** *trace-name*] [**detail**]

— **session sap** *sap-id* **mac** *ieee-address* [{**circuit-id** *circuit-id* | **remote-id** *remote-id*}] [**detail**]
— **trace** [**name** *trace-name*]
— **status**
— **trace-profile** [**detail**]
— **trace-profile** *profile-name*

9.25.1.28 Monitor Commands

monitor
— **service**
— **subscriber** *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**] [**base**] [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]
— **subscriber** *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**] **egress-policer-id** *egress-policer-id* [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]
— **subscriber** *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**] **egress-queue-id** *egress-queue-id* [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]
— **subscriber** *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**] **ingress-policer-id** *ingress-policer-id* [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]
— **subscriber** *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**] **ingress-queue-id** *ingress-queue-id* [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]

9.25.1.29 Clear Commands

clear
- **call-trace**
  - **ipoe**
    - **session sap** *sap-id* **mac** *ieee-address* [{**circuit-id** *circuit-id* | **remote-id** *remote-id*}]
    - **trace** *trace-name*
- **subscriber-mgmt**
  - **accu-stats active-subs no-accu-stats-policy**
  - **accu-stats active-subs sub-profile** *profile-name*
  - **accu-stats inactive-subs**
  - **accu-stats subscriber** *subscriber-id*
  - **ancp**
    - **ancp-sub-string** *string*
  - **authentication** [*policy-name*]
  - **authentication coa-statistics**
  - **bonding** *bonding-id*
  - **diameter-session**
    - **ccrt-replay diameter-application-policy** *name* **statistics**
    - **ccrt-replay diameter-application-policy** *name* **sessions**
  - **errors**
  - **idle-only-msap** *msap-policy-name*
  - **msap-policy** *msap-policy-name* [**idle-only**]
  - **peakvalue-stats iom** (*slot* | **all**) [**recursive**]
  - **peakvalue-stats mda** (*mda* | **all**) [**recursive**]
  - **peakvalue-stats port** (*port-id* | **all**)
  - **peakvalue-stats pw-port** (*pw-port* | **all**)
  - **peakvalue-stats system** [**recursive**]
  - **radius-accounting** [*policy-name*]
  - **statistics sla-profile** {*sla-profile-name* | **all**}
  - **statistics sub-profile** {*sub-profile-name* | **all**}
- **service**
  - **id** *service-id*
    - **ipoe session** [**sap** *sap-id*] [**interface** *ip-int-name* | *ip-address*] [**mac** *ieee-address*] [**circuit-id** *circuit-id*] [**remote-id** *remote-id*] [**inter-dest-id** *intermediate-destination-id*] [**no-inter-dest-id**] [**ip-address** *ip-prefix*[/*prefix-length*]] [**port** *port-id*] [**subscriber** *sub-ident-string*] [**sap-session-id** *sap-session-index*]
    - **ipoe session all**
    - **msap** *msap-id*
    - **msap-policy** *msap-policy-name*
  - **statistics**
    - **subscriber** *sub-ident-string*
- **qos**
  - **scheduler-stats**
    - **subscriber** *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* [**scheduler** *scheduler-name*] [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]
    - **subscriber** *sub-ident-string* [**scheduler** *scheduler-name*] [**ingress** | **egress**]

clear
- **router**

```
            — srrp
                — interface subscriber-interface [id srrp-id]
                — statistics interface subscriber-interface [id srrp-id]

clear
    — aaa
        — route-downloader name [vprn vprn] [family family]
```

9.25.1.30 Debug Commands

```
debug
    — call-trace
        — ipoe
            — trace sap sap-id [mac ieee-address] [circuit-id circuit-id] [remote-id remote-id] [profile trace-profile-name] [trace-existing-sessions] [max-jobs num] [name trace-name]
            — trace mac ieee-address [sap sap-id] [circuit-id circuit-id] [remote-id remote-id] [profile trace-profile-name] [trace-existing-sessions] [max-jobs num] [name trace-name]
            — trace circuit-id circuit-id [profile trace-profile-name] [trace-existing-sessions] [max-jobs num] [name trace-name]
            — trace remote-id remote-id [profile trace-profile-name] [trace-existing-sessions] [max-jobs num] [name trace-name]
            — no trace [sap sap-id] [mac ieee-address] [circuit-id circuit-id] [remote-id remote-id]
            — no trace name trace-name

debug
    — [no] diameter
        — application
            — [no] on-error
            — [no] policy policy-name
                — [no] on-error
                — [no] session-messages
        — dest-realm realm
        — no dest-realm
        — detail-level level
        — diameter-peer peer [psm-events]
        — no diameter-peer
        — diameter-peer-policy policy
        — no diameter-peer-policy policy
            — avp-match id
            — no avp-match id
                — avp avp-id type type [ascii | hex] value value
                — no avp avp-id
                — message-type [ccr] [cca] [rar] [raa] [asr] [asa] [aar] [aaa]
                — message-type all
                — no message-type
                — [no] shutdown
            — diameter-peer peer [psm-events]
```

```
                    — no diameter-peer
                    — message-type [ccr] [cca] [cer] [cea] [dwr] [dwa] [dpr] [dpa] [rar] [raa] [asr]
                        [asa] [aar] [aaa]
                    — message-type all
                    — no message-type
                — message-type [ccr] [cca] [cer] [cea] [dwr] [dwa] [dpr] [dpa] [rar] [raa] [asr] [asa] [aar]
                    [aaa]
                — message-type all
                — no message-type
                — [no] node host-name
                    — [no] on-error
                    — [no] peer destination-host
                        — [no] on-error
                        — [no] peer-to-peer
                — origin-realm realm
                — no origin-realm

debug
    — service
        — id service-id
            — arp
            — [no] arp-host
            — one-time-http-redirection
            — [no] ppp
                — [no] circuit-id circuit-id
                — [no] event
                    — dhcp-client [terminate-only]
                    — no dhcp-client
                    — l2tp [terminate-only]
                    — no l2tp
                    — local-address-assignment [terminate-only]
                    — no local-address-assignment
                    — ppp [terminate-only]
                    — no ppp
                — [no] mac ieee-address
                — [no] msap msap-id
                — [no] packet
                    — detail-level {low | medium | high}
                    — no detail-level
                    — [no] dhcp-client
                    — discovery [padi] [pado] [padr] [pads] [padt]
                    — no discovery
                    — mode {dropped-only | ingr-and-dropped | egr-ingr-and-
                        dropped}
                    — no mode
                    — ppp [lcp] [pap] [chap] [ipcp] [ipv6cp]
                    — no ppp
                — [no] remote-id remote-id
                — [no] sap sap-id
                — [no] username username
    — subscriber-mgmt
            — authentication [policy policy-name] [mac-addr ieee-address] [circuit-id
                circuit-id]
            — [no] sub-ident-policy policy-name
```

```
                    — script-all-info
                    — [no] script-compile-error
                    — [no] script-export-variables
                    — [no] script-output
                    — [no] script-output-on-error
                    — [no] script-runtime-error
          — router
              — [no] srrp
                  — [no] events [interface ip-int-name]
                  — [no] packets [interface ip-int-name]
              — [no] radius
                  — detail-level {low | medium | high}
                  — no detail-level
                  — packet-type [authentication] [accounting] [coa]
                  — no packet-type
                  — radius-attr type attribute-type [extended-type attribute-ext-type]
                      [transaction]
                  — radius-attr type attribute-type [extended-type attribute-ext-type]
                      [transaction] {address | hex | integer | string} value attribute-value
                  — radius-attr vendor vendor-id type attribute-type [extended-type attribute-ext-
                      type] [transaction encoding-type]
                  — radius-attr vendor vendor-id type attribute-type [extended-type attribute-ext-
                      type] [transaction encoding-type] {address | hex | integer | string} value
                      attribute-value
                  — no radius-attr type attribute-type [extended-type attribute-ext-type]
                  — no radius-attr type attribute-type [extended-type attribute-ext-type] {address
                      | hex | integer | string} attribute-value
                  — no radius-attr vendor vendor-id type attribute-type [extended-type attribute-
                      ext-type]
                  — no radius-attr vendor vendor-id type attribute-type [extended-type attribute-
                      ext-type] {address | string} value attribute-value
                  — [no] server-address ip-address
              — [no] wpp
                  — [no] packet
                      — detail-level detail-level
                  — [no] portal wpp-portal-name
                      — [no] packet
                          — detail-level detail-level
```

9.25.1.31 Tools Commands

```
tools
    — dump
        — debug
            — diameter
                — avp-match-learned-session-id
        — redundancy
            — multi-chassis
                — mc-endpoint peer ip-address
                — mc-ring
                — mc-ring peer ip-address [ring sync-tag]
```

— srrp-sync-database [instance instance-id] [peer ip-address]
— sync-database [peer *ip-address*] [port *port-id* | *lag-id*] [sync-tag *sync-tag*] [application *application*] [detail] [type *type*]
— sync-database [peer *ip-address*] [sdp *sdp-id*] [sync-tag *sync-tag*] [application *application*] [detail] [type *type*]

— perform
— aaa
— route-downloader start [force]
— redundancy
— multi-chassis
— multi-chassis
— mc-ipsec
— force-switchover tunnel-group *local-group-id*
— sync-database-reconcile [peer *ip-address*] [sdp *sdp-id* [sync-tag *sync-tag*]] [application *application*]
— sync-database-reconcile [peer *ip-address*] [port *port-id* | *lag-id* [sync-tag *sync-tag*]] [application *application*]
— persistence
— downgrade target-version *target* [reboot]
— subscriber-mgmt
— bonding
— eval-rates
— coa [nas-port-id *attr-string*] [framed-ip-addr *ip-address*] [alc-ipv6-addr *ipv6-address*] [delegated-ipv6-pfx *ipv6-prefix*[/*ipv6-prefix-length*] [framed-ipv6-pfx *ipv6-prefix*[/*ipv6-prefix-length*]] [alc-serv-id *serv-id*] [acct-session-id *attr-string*] [alc-subscr-id *attr-string*] [alc-brg-id *attr-string*] [alc-client-hw-addr *mac-address*] [attr *attribute* [*attribute*...(up to 5 max)]] [attr-from-file *file-url*] [from-server *attr-string*] [router-or-service *attr-string*] [debug]
— credit-reset sap *sap-id* subscriber *sub-ident-string* sla-profile *sla-profile-name* {category *category-name* | all-categories} [spi-sharing-type *spi-sharing-type*] [spi-sharing-id *spi-sharing-id*
— credit-reset sap *sap-id* ip *ip-address* {category *category-name* | all-categories} [mac *mac-address*]
— credit-reset svc *service-id* ip *ip-address* {category *category-name* | all-categories} [mac *mac-address*]
— edit-ipoe-session sap *sap-id* mac *mac-address* [subscriber *sub-ident-string*] [sub-profile-string *sub-profile-string*] [sla-profile-string sla-*profile-string*] [inter-dest-id *intermediate-destination-id*] [ancp-string *ancp-string*] [app-profile-string *app-profile-string*] [circuit-id *circuit-id*] [remote-id *remote-id*]
— edit-lease-state sap *sap-id* ip *ip-address* [mac *ieee-address*] [subscriber *sub-ident-string*] [sub-profile-string *sub-profile-string*] [sla-profile-string *sla-profile-string*] [app-profile-string *app-profile-string*] [inter-dest-id *intermediate-destination-id*] [ancp-string *ancp-string*]
— edit-lease-state svd-id *service-id* ip *ip-address* [mac *ieee-address*] [subscriber *sub-ident-string*] [sub-profile-string *sub-profile-string*] [sla-profile-string *sla-profile-string*] [app-profile-string *app-profile-string*] [inter-dest-id *intermediate-destination-id*] [ancp-string *ancp-string*]
— edit-ppp-session sap *sap-id* [user-name *user-name*] [mac *mac-address*] [session-id *session-id*] [ip-address *ip-address*] [subscriber *sub-ident-string*] [sub-profile-string *sub-profile-string*] [sla-profile-string *sla-profile-string*] [inter-dest-id *intermediate-destination-id*] [ancp-string *ancp-string*] [app-profile-string *app-profile-string*]

— **edit-ppp-session** **svc-id** *service-id* [**user-name** *user-name*] [**mac** *mac-address*] [**session-id** *session-id*] [**ip-address** *ip-address*] [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*] [**inter-dest-id** *intermediate-destination-id*] [**ancp-string** *ancp-string*] [**app-profile-string** *app-profile-string*]
— **edit-slaac-host** **sap** *sap-id* **ipv6-address** *ipv6-prefix* [**mac** *ieee-address*] [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*] [**app-profile-string** *app-profile-string*] [**inter-dest-id** *intermediate-destination-id*] [**ancp-string** *ancp-string*]
— **edit-slaac-host** **svc-id** *service-id ipv6-prefix* [**mac** *ieee-address*] [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*] [**app-profile-string** *app-profile-string*] [**inter-dest-id** *intermediate-destination-id*] [**ancp-string** *ancp-string*]
— **eval-ipoe-session** [**svc-id** *service-id*] [**sap** *sap-id*] [**mac** *mac-address*] [**circuit-id** *circuit-id*] [**remote-id** *remote-id*] [**subscriber** *sub-ident-string*]
— **eval-ipoe-session** [**svc-id** *service-id*] [**sap** *sap-id*] [**mac** *mac-address*] [**circuit-id** *circuit-id*] [**remote-id** *remote-id*] [**subscriber** *sub-ident-string*]
— **eval-lease-state** **sap** *sap-id* **ip** *ip-address* [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*] [**app-profile-string** *app-profile-string*] [**inter-dest-id** *intermediate-destination-id*] [**ancp-string** *ancp-string*]
— **eval-lease-state** **svc-id** *service-id* **ip** *ip-address* [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*] [**app-profile-string** *app-profile-string*] [**inter-dest-id** *intermediate-destination-id*] [**ancp-string** *ancp-string*]
— **eval-msap** {**policy** *msap-policy-name* | **msap** *sap-id*}
— **eval-ppp-session** [**svc-id** *service-id*] [**user-name** *user-name*] [**sap** *sap-id*] [**ip** *ip-address*] [**subscriber** *sub-ident-string*]
— **eval-ppp-session** [**svc-id** *service-id*] [**user-name** *user-name*] [**sap** *sap-id*] **mac** *mac-address* [**session-id** *session-id*] [**ip** *ip-address*] [**subscriber** *sub-ident-string*]
— **eval-slaac-host** [**svc-id** *service-id*] [**sap** *sap-id*] [**subscriber** *sub-ident-string*] [**ipv6-address** *ipv6-prefix*] [**mac** *ieee-address*]
— **forcerenew** **svc-id** *service-id* {**ip** *ip-address*[*/mask*]**>** | **mac** *ieee-address*}
— **forcerenew** {**interface** *interface-name* | **sap** *sap-id* | **sdp** *sdp-id:vc-id*} [**ip** *ip-address*[*/mask*] | **mac** *ieee-address*]
— **local-user-db** *local-user-db-name*
    — **ipoe**
        — **host-lookup** [**mac** *ieee-address*] [**remote-id** *remote-id-ascii*] [**sap-id** *sap-id*] [**service-id** *service-id*] [**string** *vso-string*] [**system-id** *system-id]* [**option60** *option-60-ascii*] [**circuit-id** *circuit-id-ascii*] [**circuit-id-hex** circuit-id-hex] [**option60-hex** *option60-hex*] [**remote-id-hex** *remote-id-hex*] [**derived-id** *derived-id*] [**ip-prefix** *ip-prefix/ip-prefix-length*]
    — **ppp**
        — **authentication** **user-name** *user-name* **password** *password* [**circuit-id** *circuit-id*] [**circuit-id-hex** *circuit-id-hex*] [**derived-id** *derived-id*] [**mac** *ieee-address*] [**remote-id** *remote-id*] [**remote-id-hex** *remote-id-hex*] [**sap-id** *sap-id*] [**service-name** *service-name*]

— **host-lookup** [**circuit-id** *circuit-id*] [**circuit-id-hex** *circuit-id-hex*] [**derived-id** derived-id] [**mac** *ieee-address*] [**remote-id** *remote-id*] [**remote-id-hex** *remote-id-hex*] [**sap-id** *sap-id*] [**service-name** *service-name*] [**user-name** *user-name*]
— **re-ident-sub** *old-sub-ident-string* **to** *new-sub-ident-string*
— **eval-msap**{**policy** *msap-policy-name* | **msap** *sap-id*}
— **remap-lease-state** **old-mac** *ieee-address* **mac** *ieee-address*
— **remap-lease-state** *sap-id* [**mac** *ieee-address*]

9.25.2 Triple Play Subscriber Management Configuration Commands

9.25.2.1 Triple Play Subscriber Management Configuration Commands

- Generic Commands
- ANCP Commands
- BGP Peering Policy Commands
- Authentication Policy Commands
- RADIUS Route Download Commands
- IGMP Policy Commands
- PIM Policy Commands
- Managed SAP Policy Commands
- Multi-Chassis Redundancy Commands
- Router Advertisement Policy Commands
- SLA Profile Commands
- Subscriber Identification Policy Commands
- Subscriber Profile Commands
- Subscriber Management Service Commands
- Subscriber Management Service Commands
- RIP Commands
- Vport Commands
- Clear Commands
- Show Commands
- Tools Commands
- Debug Commands

- Monitor Commands

9.25.2.1.1 Generic Commands

description

Syntax **description** *description-string*
no description

Context config>subscr-mgmt>accu-stats-policy
config>subscr-mgmt>authentication-policy
config>subscr-mgmt>host-tracking
config>subscr-mgmt>pim-policy
config>subscr-mgmt>sla-profile
config>subscr-mgmt>sla-profile>egress>ip-filter>entry
config>subscr-mgmt>sla-profile>ingress>ip-filter>entry
config>subscr-mgmt>sub-ident-pol
config>subscr-mgmt>sub-profile
config>subscr-mgmt>mld-policy
config>service>vpls>gsmp>group
config>log>accounting-policy
config>service>vprn>redundant-interface
config>service>ies>redundant-interface
config>service>ies>subscriber-interface
config>service>ies>sub-if>grp-if
config>service>ies>sub-if>grp-if>srrp
config>service>vprn>subscriber-interface
config>service>vprn>sub-if>grp-if
config>service>vprn>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>srrp
config>service>vprn>gsmp>group
config>service>vprn>gsmp>group>neighbor
config>service>vprn>sub-if>grp-if>ipoe-session
config>redundancy>multi-chassis>peer
config>subscr-mgmt>cat-map
config>subscr-mgmt>ipoe-session-policy
config>service>vpls>sap> ipoe-session
config>subscr-mgmt>diam-appl-plcy
config>sub-mgmt>credit-control-policy
config>sub-mgmt>host-lockout>policy
config>subscr-mgmt>sub-mcac-policy
config>aaa>route-downloader
config>aaa>diam-peer-plcy
config>aaa>wpp>portal-groups
config>aaa>wpp>portal-groups>portal-group
config>port>ethernet>access>egress

config>subscr-mgmt>cat-map>category
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6
config>call-trace>trace-profile
config>aaa>diam>node

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes any description string from the context.

Parameters *description-string* — Specifies a text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

shutdown

Syntax [**no**] **shutdown**

Context config>subscr-mgmt>sub-ident-pol>primary
config>subscr-mgmt>sub-ident-pol>secondary
config>subscr-mgmt>sub-ident-pol>tertiary
config>service>vpls>sap>sub-sla-mgmt
config>service>vpls>gsmp
config>service>vpls>gsmp>group
config>service>vpls>gsmp>group>neighbor
config>service>vprn>redundant-interface
config>service>vprn>red-if>spoke-sdp
config>service>vprn>subscriber-interface
config>service>vprn>sub-if>grp-if
config>service>vprn>sub-if>grp-if>srrp
config>service>ies>subscriber-interface
config>service>ies>sub-if>grp-if>srrp
config>service>ies>redundant-interface
config>service>ies>sub-if>grp-if>arp-host
config>service>vprn>gsmp>group>neighbor
config>service>ies>sub-if>grp-if>wpp
config>service>vprn>sub-if>grp-if>wpp
config>service>vprn>sub-if>grp-if>wpp>portals
config>redundancy>multi-chassis>peer
config>redundancy>multi-chassis>peer>mc-lag
config>redundancy>multi-chassis>peer>sync
config>subscr-mgmt>sub-mcac-policy
config>aaa>route-downloader
config>aaa>diam-peer-plcy>peer

config>aaa>diam-peer-plcy>proxy
config>aaa>wpp>portal-groups
config>aaa>wpp>portal-groups>portal-group
config>service>vprn>sub-if>grp-if>data-trigger
config>service>ies>sub-if>grp-if>data-trigger
config>service>ies>sub-if>grp-if>gtp-parameters
config>service>vprn>sub-if>grp-if>gtp-parameters
config>service>ies>sub-if>grp-if>srrp
config>service>vprn>sub-if>grp-if>srrp
config>aaa>diam>node>peer

Description The **shutdown** command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command.

The **shutdown** command administratively disables an entity. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they can be deleted.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system generated configuration files.

Shutting down a subscriber interface on a 7750 SR will operationally shut down all child group interfaces and SAPs. Shutting down a group interface on a 7750 SR will operationally shut down all SAPs that are part of that group-interface.

The **no** form of this command puts an entity into the administratively enabled state.

ies

Syntax **ies** *service-id* [**customer** *customer-id*] [**create**] [**vpn** *vpn-id*] [**name** *name*]
no ies *service-id*

Context config>service

Description This command creates or edits an IES service instance.

The **ies** command creates or maintains an Internet Ethernet Service (IES). If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

IES services allow the creation of customer facing IP interfaces in the same routing instance used for service network core routing connectivity. IES services require that the IP addressing scheme used by the subscriber must be unique between it and other addressing schemes used by the provider and potentially the entire Internet.

While IES is part of the routing domain, the usable IP address space may be limited. This allows a portion of the service provider address space to be set aside for service IP provisioning, becoming administered by a separate but subordinate address authority. This feature is defined using the **config router service-prefix** command.

IP interfaces defined within the context of an IES service ID must have a SAP created as the access point to the subscriber network. This allows a combination of bridging and IP routing for redundancy purposes.

When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the **customer** command in the service context. Once a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and recreated with a new customer association.

Once a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified results in an error.

Multiple IES services are created to separate customer owned IP interfaces. More than one IES service may be created for a single customer ID. More than one IP interface may be created within a single IES service ID. All IP interfaces created within an IES service ID belongs to the same customer.

By default, no IES service instances exist until they are explicitly created.

The **no** form of this command deletes the IES service instance with the specified *service-id*. The service cannot be deleted until all the IP interfaces defined within the service ID have been shutdown and deleted.

Parameters *service-id* — Specifies the unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every router on which this service is defined.

Values *service-id*: 1 to 214748364
svc-name: A string up to 64 characters

customer-id — Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn-id — Specifies the VPN ID number used to identify virtual private networks (VPNs) by a VPN identification number.

Values 1 to 2147483647

Default **null** (0)

create — Keyword used to create the service ID. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

vpls

Syntax **vpls** *service-id* [**customer** *customer-id*] [*create*] **vpn** *vpn-id* [**m-vpls**] [**bvpls** | **i-vpls**] [**etree**] [**create**] [**name** *name*]
no vpls *service-id*

Context config>service

Description This command enables the context to configure a Virtual Private LAN Services (VPLS) instance. The **vpls** command is used to create or maintain a VPLS service. If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

A VPLS service connects multiple customer sites together acting like a zero-hop, layer 2 switched domain. A VPLS is always a logical full mesh.

When a service is created, the **create** keyword must be specified if the **create** command is enabled in the **environment** context.

When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the **customer** command in the service context. Once a service has been created with a customer association, it is not possible to edit the customer association. The service must be deleted and recreated with a new customer association.

To create a management VPLS, the **m-vpls** keyword must be specified. See section **Hierarchical VPLS Redundancy** for an introduction to the concept of management VPLS.

Once a service is created, the use of the **customer** *customer-id* is optional for navigating into the service configuration context. Attempting to edit a service with the incorrect *customer-id* specified results in an error.

More than one VPLS service may be created for a single customer ID.

By default, no VPLS instances exist until they are explicitly created.

The **no** form of this command deletes the VPLS service instance with the specified *service-id*. The service cannot be deleted until all SAPs and SDPs defined within the service ID have been shutdown and deleted, and the service has been shutdown.

Parameters *service-id* — Specifies the unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every router on which this service is defined.

Values *service-id*: 1 to 214748364
svc-name: A string up to 64 characters.

customer-id — Specifies the customer ID number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

vpn-id — Specifies the VPN ID number which allows you to identify virtual private networks (VPNs) by a VPN identification number

Values 1 to 2147483647

Default **null** (0)

m-vpls — Specifies a managed VPLS

create — Keyword used to create a VPLS service. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

vprn

Syntax **vprn** *service-id* [**customer** *customer-id*] [**create**]
no vprn *service-id*

Context config>service

Description This command creates or edits a Virtual Private Routed Network (VPRN) service instance.

If the *service-id* does not exist, a context for the service is created. If the *service-id* exists, the context for editing the service is entered.

VPRN services allow the creation of customer-facing IP interfaces in the same routing instance used for service network core routing connectivity. VPRN services require that the IP addressing scheme used by the subscriber must be unique between it and other addressing schemes used by the provider and potentially the entire Internet.

IP interfaces defined within the context of an VPRN service ID must have a SAP created as the access point to the subscriber network.

When a service is created, the **customer** keyword and *customer-id* must be specified and associates the service with a customer. The *customer-id* must already exist having been created using the customer command in the service context. When a service is created with a customer association, it is not possible to edit the customer association. The service must be deleted and re-created with a new customer association.

When a service is created, the use of the **customer** *customer-id* is optional to navigate into the service configuration context. If attempting to edit a service with the incorrect *customer-id* results in an error.

Multiple VPRN services are created to separate customer-owned IP interfaces. More than one VPRN service can be created for a single customer ID. More than one IP interface can be created within a single VPRN service ID. All IP interfaces created within an VPRN service ID belongs to the same customer.

The **no** form of this command deletes the VPRN service instance with the specified *service-id*. The service cannot be deleted until all the IP interfaces and all routing protocol configurations defined within the service ID have been shutdown and deleted.

Parameters *service-id* — Specifies the unique service identification number identifying the service in the service domain. This ID must be unique to this service and may not be used for any other service of any type. The *service-id* must be the same number used for every 7750 SR on which this service is defined.

Values

service-id:	1 to 2147483648
svc-name:	64 characters maximum

customer-id — Specifies an existing customer identification number to be associated with the service. This parameter is required on service creation and optional for service editing or deleting.

Values 1 to 2147483647

create — Keyword used to create a VPRN service. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

subscriber-mgmt

Syntax **subscriber-mgmt**

Context config

Description This command enables the context to configure subscriber management entities. A subscriber is uniquely identified by a subscriber identification string. Each subscriber can have several DHCP sessions active at any time. Each session is referred to as a subscriber host and is identified by its IP address and MAC address.

All subscriber hosts belonging to the same subscriber are subject to the same hierarchical QoS (HQoS) processing. The HQoS processing is defined in the sub-profile (the subscriber profile). A sub-profile refers to an existing scheduler policy (configured in **the configure>qos>scheduler-policy** context) and offers the possibility to overrule the rate of individual schedulers within this policy.

Because all subscriber hosts use the same scheduler policy instance, they must all reside on the same complex.

9.25.2.1.2 ANCP Commands

ancp

Syntax **ancp**

Context config>subscr-mgmt
config>subscr-mgmt>sub-prof

Description This command enables the context to configure Access Node Control Protocol (ANCP) parameters.

ancp-policy

Syntax **ancp-policy** *policy-name* [**create**]
no ancp-policy *policy-name*

Context config>subscr-mgmt>ancp

Description This command creates an Access Node Control Protocol (ANCP) policy. The policy is associated with either the ANCP string (static case) or subscriber-profile (dynamic case) and defines the behavior of the hosts belonging to these profiles.

ANCP polices control rates and subscribers based on port-up/port-down messages from the access node. When configured, the 7450 ESS or 7750 SR should stop SHCV to a host that is part of a port defined to be down (by port-down message). When the node receives a port-up message for a port that was in port-down, state the node will initiate the SHCV process immediately to verify connectivity.

When ANCP is used with Enhanced Subscriber Management, the ANCP string last associated with the subscriber is used. All hosts of a subscriber is updated with the new ANCP string.

The **no** form of this command removes the policy name from the ANCP configuration.

Parameters *policy-name* — Configures the ANCP policy name, up to 32 characters

create — Keyword used to create the ANCP policy. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

egress

Syntax **egress**

Context config>subscr-mgmt>ancp>ancp-policy

Description This command enables the context to configure egress ANCP policy parameters.

rate-adjustment

Syntax **rate-adjustment** *adjusted-percent*
no rate-adjustment

Context config>subscr-mgmt>ancp>ancp-policy>ingress
config>subscr-mgmt>ancp>ancp-policy>egress

Description This command configures a rate adjustment for the scheduler. The **rate-adjustment** command should be used when the rate returned by the DSLAM is calculated with different encapsulation than the 7450 ESS or 7750 SR. The node will adjust the rate by the percent specified as:

DSLAM_RATE*adjust-rate/100 — rate-reduction.

The **no** form of this command reverts to the default.

Parameters *adjusted-percent* — Specifies a rate adjustment for the scheduler

Values 1 to 200

Default 100

rate-modify

Syntax **rate-modify scheduler** *scheduler-name*
rate-modify agg-rate-limit
no rate-modify

Context config>subscr-mgmt>ancp>ancp-policy>egress
config>subscr-mgmt>ancp>ancp-policy>ingress

Description This command configures rate modify scheduler parameters.

The **no** form of this command removes the scheduler name from the configuration.

Parameters **agg-rate-limit** — Specifies that the maximum total rate for all subscriber egress or ingress queues for each subscriber associated with the policy.

scheduler-name — Specifies a scheduler name up to 32 characters

rate-monitor

Syntax **rate-monitor** *kilobit-per-second* [**alarm**]
no rate-monitor

Context config>subscr-mgmt>ancp>ancp-policy>ingress
config>subscr-mgmt>ancp>ancp-policy>egress

Description This command configures the rate monitor level.

The **no** form of this command removes the value from the configuration.

Parameters *kilobit-per-second* — Specifies the rate, in kilobits, below which the system will generate an event

alarm — When the monitored rate is below the configured value the system generates an alarm (trap) to the management system. The trap includes the rate as well as the ANCP policy name and the ANCP string.

rate-reduction

Syntax **rate-reduction** *kilobit-per-second*
no rate-reduction

Context config>subscr-mgmt>ancp>ancp-policy>ingress
config>subscr-mgmt>ancp>ancp-policy>egress

Description This command defines a constant rate reduction to the rate specified by the DSLAM. The **rate-reduction** command should be used if the node should adjust the rate to a value that is offset (for example by a fixed multicast dedicated bandwidth) compared to the total available on the DSLAM.

When set, the rate is:

DSLAM_RATE*adjust-rate/100 — rate-reduction

The **no** form of this command removes the value from the configuration.

ingress

Syntax **ingress**

Context config>subscr-mgmt>sla-prof>ingress
config>subscr-mgmt>ancp>ancp-policy

Description This command enables the context to configure ingress ANCP policy parameters.

port-down

Syntax [**no**] **port-down**

Context config>subscr-mgmt>ancp>ancp-policy

Description This command enables the context to configure the actions taken on port-down.

disable-shcv

Syntax [**no**] **disable-shcv** [**alarm**] [**hold-time** *seconds*]

Context config>subscr-mgmt>ancp>ancp-policy>port-down

Description When this command is configured, the node will suspend SHCV for the hosts defined with this ANCP policy until the access node sends a port-up message. When the hold-time parameter is used, the node will suspend SHCV for the period of time defined. If the hold-time parameter is not defined the node will suspend SHCV until a port-up message is received.

If the optional alarm flag is used the node should send a SHCV alarm before suspending.

The **no** form of this command reverts to the default.

ancp-static-map

Syntax **ancp-static-map**

Context config>subscr-mgmt>ancp

Description This command enables the context to configure a static ANCP name map.

Default ancp-static-map

entry

Syntax **entry key** *ancp-string* **customer** *customer-id* **multi-service-site** *customer-site-name* **ancp-policy** *policy-name*
entry key *ancp-string* **sap** *sap-id* **ancp-policy** *policy-name*
no entry key *ancp-string* **customer** *customer-id* **multi-service-site** *customer-site-name*
no entry key *ancp-string* **sap** *sap-id*

Context config>subscr-mgmt>ancp>static-map

Description This command configures an ANCP name. When ANCP is configured to provide rate adaptation without the use of enhanced subscriber management, this command will define how to map an ANCP key (usually the circuit-id of the DSLAM port) to either a SAP and a scheduler name (when a Multi-Service Site (MSS) is not used) or a customer, site and scheduler name when MSS is used.

Different ANCP names may be used with the same SAPs or customer ID/MSS combinations to allow schedulers within the policy to be mapped to the ANCP names. An ANCP string and SAP combination may reference only one ancp-policy. An ANCP string and customer and site-name combination may reference a single ancp-policy.

The **no** form of this command reverts to the default.

Parameters *ancp-string* — Specifies the ASCII representation of the DSLAM circuit-id name up to 63 characters

customer-id — Specifies the associated existing customer ID

Values 1 to 2147483647

customer-site-name — Specifies the associated customer's configured MSS name, up to 32 characters

policy-name — Specifies an existing ANCP policy name up to 32 characters

sap-id — Specifies the physical port identifier portion of the SAP definition

9.25.2.1.3 Authentication Policy Commands

radius-coa-port

Syntax **radius-coa-port** {*port-number*}
no radius-coa-port

Context config>aaa

Description This command configures the system-wide UDP port number that RADIUS is listening on for CoA and Disconnect messages

The **no** form of this command reverts to the default.

Default radius-coa-port 3799

Parameters *port-number* — Specifies the UDP port number for RADIUS CoA and disconnect messages

Values 1647, 1700, 1812, 3799

authentication-policy

Syntax **authentication-policy** *name* [**create**]
no authentication-policy

Context config>subscr-mgmt
config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command creates the context to configure RADIUS server parameters for session authentication. The policies can be applied to an IES or VPRN interface, or a VPLS SAP.

The **no** form of this command removes the RADIUS server configuration for session authentication.

RADIUS servers can be configured for three different applications:

1. For authentication of dynamic Triple Play subscriber sessions, under **config>subscr-mgmt>authentication-plcy**.

2. For 802.1x port authentication, under **config>system>security>dot1x>radius-plcy**.
3. For CLI login users, under **config>system>radius**.

Parameters *name* — Specifies the name of the profile. The string is case sensitive and limited to 32 ASCII 7-bit printable characters with no spaces.

create — Keyword used to create the authentication policy. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

accept-authorization-change

Syntax [**no**] **accept-authorization-change**

Context config>subscr-mgmt>auth-policy

Description This command specifies whether or not the system should handle the CoA messages initiated by the RADIUS server, and provide for mid-session interval changes of policies applicable to subscriber hosts.

The **no** form of this command reverts to the default.

accept-script-policy

Syntax **accept-script-policy** *policy-name*
no accept-script-policy

Context config>subscr-mgmt>auth-policy

Description This command configures a RADIUS script policy used to change the RADIUS attributes of the incoming Access-Accept messages.

The **no** form of this command reverts to the default.

Parameters *policy-name* — Specifies a Python script policy to modify Access-Accept messages up to 32 characters

coa-script-policy

Syntax **coa-script-policy** *policy-name*
no coa-script-policy

Context config>subscr-mgmt>auth-policy

Description This command configures a RADIUS script to modify a change of authorization.

Parameters *policy-name* — Specifies the name of a script policy, up to 32 characters

fallback-action

Syntax **fallback-action accept** [**force-probing**]
fallback-action user-db *local-user-db-name* [**force-probing**]
no fallback-action

Context config>subscr-mgmt>auth-plcy

Description This command configures the action when no RADIUS server is available; servers are either out of service or are in a probing state.

The **no** form of this command removes the action from the configuration.

Parameters **accept** — specifies that all authentication requests are automatically accepted

local-user-db-name — specifies that the LUDB is used to authenticate to the server

force-probing — specifies that a subscriber or a test user **auth-request,** depending which arrives first, is used to probe a RADIUS server before continuing with **accept** or **user-db**. This forces an out-of-service server to transition to a probing state.

gtp-user-name

Syntax **gtp-user-name** {**imsi** | **imsi-apn**}
no gtp-user-name

Context config>subscr-mgmt>auth-plcy

Description This command defines the user-name used for authentication of a GTP session.

The **no** form of this command reverts to the default.

Default gtp-user-name imsi

Parameters **imsi** — Specifies that the user-name only contains the imsi

imsi-apn — Specifies that the user-name only contains the imsi and APN

access-loop-options

Syntax [**no**] **access-loop-options**

Context config>subscr-mgmt>auth-plcy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables inclusion of access loop information: Broadband Forum (BBF) access loop characteristics, DSL line state and DSL type. The BBF access loop characteristics are returned as BBF specific RADIUS attributes where DSL line state and DSL type are returned as Nokia-specific RADIUS VSAs.

Information obtained via the ANCP protocol has precedence over information received in PPPoE Vendor Specific BBF tags or DHCP Vendor Specific BBF Options.

If ANCP is utilized and interim accounting update is enabled, any Port Up event from GSMP will initiate in an interim update. Port Up messages can include information such as an update on the current subscriber actual-upstream-speed. The next interim accounting message is from port up triggering point.

The **no** form of this command reverts to the default.

host-accounting

Syntax [**no**] **host-accounting** [**interim-update**]

Context config>subscr-mgmt>acct-plcy

Description This command enables per-host accounting. In host accounting mode, the acct-session-id is generated per host. This acct-session-id is uniformly included in all accounting messages (START/INTERIM-UPDATE/STOP) and it can be included in RADIUS Access-Request message.

Accounting counters are based on the queue counters and as such are aggregated for all host sharing the queues within an sla-profile instance (non HSMDA) or a subscriber (HSMDA). CoA and LI is supported based on the acct-session-id of the host.

The **no** form of this command reverts to the default.

Parameters **interim-update** — Specifies that no interim-update messages are sent for the related subscriber hosts when the session is deleted. Without this keyword, only START and STOP accounting messages are generated when the host is established or terminated. This is equivalent to a time-based accounting where only the duration of the session is required.

include-radius-attribute

Syntax [**no**] **include-radius-attribute**

Context config>subscr-mgmt>auth-plcy
config>subscr-mgmt>acct-plcy

Description This command enables the context to specify the RADIUS parameters that the system should include into RADIUS authentication-request messages.

acct-authentic

Syntax [**no**] **acct-authentic**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the acct-authentic RADIUS attribute and only applies to the 7750 SR.

The **no** form of this command reverts to the default.

apn

Syntax [**no**] **apn**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute
config>subscr-mgmt>auth-plcy>include-radius-attribute

Description This command enables the inclusion of the APN n AAA protocols as signaled in the incoming GTP setup message.

The **no** form of this command disables the inclusion of the attribute.

called-station-id

Syntax [**no**] **called-station-id**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command includes called station id attributes and only applies to the 7750 SR.

The **no** form of this command reverts to the default.

calling-station-id

Syntax **calling-station-id**
calling-station-id {**llid** | **mac** | **remote-id** | **sap-id** | **sap-string**}
no calling-station-id

Context config>service>ies>if>sap
config>service>ies>sub-if>grp-if>sap
config>service>vpls>sap
config>service>vprn>if>sap
config>service>vprn>sub-if>grp-if>sap
config>subscr-mgmt>auth-plcy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include>include-radius-attribute

Description This command enables the inclusion of the calling-station-id attribute in RADIUS authentication requests and RADIUS accounting messages. The value inserted is set at the SAP level. If no **calling-station-id** value is set at the SAP level, the **calling-station-id** attribute will not be sent.

The **no** form of this command reverts to the default.

Default calling-station-id sap-string

Parameters **illid** — Specifies that the logical link identifier (LLID) is mapping from a physical to logical identification of a subscriber line and supplied by a RADIUS llid-server

mac — Specifies that the MAC address is sent

remote-id — Specifies that the remote ID is sent

sap-id — Specifies that the SAP ID is sent

sap-string — Specifies that the value is the inserted value set at the SAP level. If no **calling-station-id** value is set at the SAP level, the **calling-station-id** attribute will not be sent.

circuit-id

Syntax [**no**] **circuit-id**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the agent-circuit-id for RADIUS.

The **no** form of this command reverts to the default.

acct-delay-time

Syntax [**no**] **acct-delay-time**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the acct-delay-time RADIUS attribute and only applies to the 7750 SR.

The **no** form of this command reverts to the default.

alc-acct-triggered-reason

Syntax [**no**] **alc-acct-triggered-reason**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command includes the alc-acct-triggered-reason attribute.

dhcp-options

Syntax [**no**] **dhcp-options**

Context config>subscr-mgmt>auth-plcy>include-radius-attribute

Description This command enables insertion of RADIUS VSA containing all DHCP options from DHCP discover (or DHCP request) message. The VSA contains all DHCP options in a form of the string. If required (the total length of all DHCP options exceeds 255B), multiple VSAs are included.

The **no** form of this command reverts to the default.

dhcp-vendor-class-id

Syntax [**no**] **dhcp-vendor-class-id**

Context config>subscr-mgmt>auth-plcy>include-radius-attribute

Description This command includes the [26-6527-36] Alc-DHCP-Vendor-Class-Id attribute in RADIUS accounting messages. The content of the DHCP Vendor-Class-Identifier option (60) is mapped in this attribute.

The **no** form of this command reverts to the default.

dhcp6-options

Syntax [**no**] **dhcp6-options**

Context config>subscr-mgmt>auth-policy>include

Description This command copies DHCPv6 options from received DHCPv6 messages on ingress access and pass them to the RADIUS server in Accept-Request. The messages is carried in the ALU VSA Alc-ToServer-Dhcp6-Options.

firewall-info

Syntax [**no**] **firewall-info**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables inclusion of the Firewall Information VSA in AAA protocols.

The **no** form of this command reverts to the default.

alc-error-code

Syntax [**no**] **alc-error-code**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables RADIUS accounting messages to include an error number and error code when the subscriber host session terminates. To obtain a complete list of error numbers and their corresponding codes, use the **tools>dump>aaa>radius-acct-terminate-cause** command.

The **no** form of this command reverts to the default.

all-authorized-session-addresses

Syntax [**no**] **all-authorized-session-addresses**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command specifies to include all included and authorized address/prefix attributes in session accounting and is applicable only for session-accounting mode.

With this flag enabled, all IP address attributes explicitly enabled to be included are the following:

- delegated-ipv6-prefix
- framed-ip-address
- framed-ip-netmask
- framed-ipv6-prefix
- ipv6-address

These are included if the corresponding addresses or prefixes are authorized (via access-accept or ludb) and independent if they are used or not.

The **no** form of this command reverts to the default.

gprs-negotiated-qos-profile

Syntax [**no**] **gprs-negotiated-qos-profile**

Context config>subscr-mgmt>auth-plcy>include-radius-attribute

Description This command enables the inclusion of the 3GPP QoS specification in AAA protocols as signaled in the incoming GTP setup message.

The **no** form of this command disables the inclusion of the attribute.

imei

Syntax [**no**] **imei**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute
config>subscr-mgmt>auth-plcy>include-radius-attribute

Description This command enables the inclusion of the IMEI in AA protocols as signaled in the incoming GTP setup message.

The **no** form of this command disables the inclusion of the attribute.

mac-address

Syntax [**no**] **mac-address**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the client MAC address RADIUS attribute.

The **no** form of this command reverts to the default.

msisdn

Syntax [**no**] **msisdn**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute
config>subscr-mgmt>auth-plcy>include-radius-attribute

Description This command enables the inclusion of the MSISDN in AAA protocols as signaled in the incoming GTP setup message.

The **no** form of this command disables the inclusion of the attribute.

nas-identifier

Syntax [**no**] **nas-identifier**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the nas-identifier RADIUS attribute.

The **no** form of this command reverts to the default.

nas-port-id

Syntax **nas-port-id**
nas-port-id [**prefix-string** *string*] [**suffix** *suffix-option*]
no nas-port-id

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the nas-port-id RADIUS attribute. Optionally, the value of this attribute (the SAP-id) can be prefixed by a fixed string and suffixed by the circuit-id or the remote-id of the client connection. If a suffix is configured, but no corresponding data is available, the suffix used is 0/0/0/0/0/0.

The **no** form of this command reverts to the default.

Parameters *string* — Specifies that a user configurable string is added to the RADIUS NAS port attribute, up to 8 characters

suffix-option — Specifies the suffix type to be added to the RADIUS NAS port attribute

Values circuit-id, remote-id

nas-port-type

Syntax **nas-port-type**
nas-port-type [*value*]
no nas-port-type

Context config>subscr-mgmt>auth-plcy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the nas-port-type RADIUS attribute. If set to **nas-port-type**, the following is sent: values: 32 (null-encap), 33 (dot1q), 34 (qinq), 15 (DHCP hosts). The **nas-port-type** can also be set as a specified value, with an integer from 0 to 255.

The **no** form of this command reverts to the default.

Parameters *value* — Specifies an enumerated integer that specifies the value that is put in the RADIUS nas-port-type attribute

Values 0 to 255

pppoe-service-name

Syntax [**no**] **pppoe-service-name**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the pppoe-service-name RADIUS attribute.

The **no** form of this command reverts to the default.

rat-type

Syntax [**no**] **rat-type**

Context config>subscr-mgmt>auth-plcy>include-radius-attribute

Description This command enables the inclusion of the Radio Access Type in AAA protocols as signaled in the incoming GTP setup message.

The **no** form of this command disables the inclusion of the attribute.

uli

Syntax [**no**] **uli**

Context config>subscr-mgmt>auth-plcy>include-radius-attribute

Description This command enables the inclusion of the User Location Information in AAA protocols as signaled in the incoming GTP setup message.

The **no** form of this command disables the inclusion of the attribute.

bonding-active-connections

Syntax [**no**] **bonding-active-connections**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command triggers the inclusion of the Alc-Bonding-Active-Connection VSA in accounting for bonding subscribers.

The **no** form of this command disables the inclusion of the attribute.

bearer-fteid

Syntax [**no**] **bearer-fteid**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command, in the case of ESM over GTP access includes the Alc-Bearer-Fteid VSA in accounting. This VSA contains the fully qualified TEID of the current GTP-U tunnel, including the bearer ID, endpoint IP addresses and TEIDs for both local and remote endpoints.

The **no** version of this command disables inclusion of the Alc-Bearer-Fteid VSA.

Default no bearer-fteid

bonding-id

Syntax [**no**] **bonding-id**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command triggers the inclusion of the Alc-Bonding-Id VSA in accounting for bonding subscribers.

The **no** form of this command disables the inclusion of the attribute.

access-loop-options

Syntax [**no**] **access-loop-options**

Context config>subscr-mgmt>auth-plcy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables inclusion of access loop information: Broadband Forum (BBF) access loop characteristics, DSL line state and DSL type. The BBF access loop characteristics are returned as BBF specific RADIUS attributes where DSL line state and DSL type are returned as Nokia-specific RADIUS VSA's.

Information obtained via the ANCP protocol has precedence over information received in PPPoE Vendor Specific BBF tags or DHCP Vendor Specific BBF Options.

The **no** form of this command reverts to the default.

acct-session-id

Syntax [**no**] **acct-session-id**

Context config>subscr-mgmt>auth-plcy>include-radius-attribute

Description The **acct-session-id** attribute for each subscriber host is generated at the very beginning of the session initiation. This command will enable or disable sending this attribute to the RADIUS server in the Access-Request messages regardless of whether the accounting is enabled or not. The **acct-session-id** attribute can be used to address the subscriber hosts from the RADIUS server in the CoA Request.

The acct-session-id attribute is unique per subscriber host network wide. It is a 22bytes long field comprised of the system MAC address along with the creation time and a sequence number in a hex format.

The **no** form of this command reverts to the default.

delegated-ipv6-prefix

Syntax [**no**] **delegated-ipv6**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the delegated-ipv6-prefix RADIUS attribute and only applies to the 7750 SR.

The **no** form of this command reverts to the default.

detailed-acct-attributes

Syntax [**no**] **detailed-acct-attributes**

Context config>subscr-mgmt>auth-plcy>include-radius-attribute

Description This command enables detailed reporting of per queue and per policer octet and packet counters using RADIUS VSAs. Enabled by default. It can be enabled simultaneously with aggregate counters (std-acct-attributes).

The **no** form of this command excludes the detailed counter VSAs from the RADIUS accounting messages.

Default detailed-acct-attributes

framed-interface-id

Syntax [**no**] **framed-interface-id**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the framed-interface-id RADIUS attribute and only applies to the 7750 SR.

The **no** form of this command reverts to the default.

framed-ip-addr

Syntax [**no**] **framed-ip-addr**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the inclusion of the framed-ip-addr attribute.

The **no** form of this command reverts to the default.

framed-ip-netmask

Syntax [**no**] **framed-ip-netmask**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the inclusion of the framed-ip-netmask attribute.

The **no** form of this command reverts to the default.

framed-ipv6-prefix

Syntax [**no**] **framed-ipv6-prefix**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the framed-ipv6-prefix RADIUS attribute and only applies to the 7750 SR.

The **no** form of this command reverts to the default.

framed-ipv6-route

Syntax [**no**] **framed-ipv6-route**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description When enabled, all valid [99] Framed-IPv6-Route attributes as received in the RADIUS authentication phase and associated with an instantiated IPv6 wan host is included in the RADIUS accounting request messages. The state of the Framed-IPv6-Route (installed, shadowed, hostInactive, and so on) is not taken into account for reporting in the accounting request messages.

The **no** form of this command reverts to the default.

framed-route

Syntax [**no**] **framed-route**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description When enabled, all valid [22] Framed-Route attributes as received in the RADIUS authentication phase and associated with an instantiated IPv4 host is included in the RADIUS accounting request messages. The state of the Framed-Route (installed, shadowed, hostInactive, and so on) is not taken into account for reporting in the accounting request messages.

The **no** form of this command reverts to the default.

ipv6-address

Syntax [**no**] **framed-ipv6-address**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the ipv6-address RADIUS attribute.

The **no** form of this command reverts to the default.

nas-port

Syntax **nas-port** *binary-spec*
no nas-port

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the nas-port RADIUS attribute. You enter decimal representation of a 32-bit string that indicates your port information. This 32-bit string can be compiled based on different information from the port (data types). By using syntax number-of-bits data-type you indicate how many bits from the 32 bits are used for the specific data type. These data types can be combined up to 32 bits in total. In between the different data types 0's and/or 1's as bits can be added.

The **no** form of this command disables your nas-port configuration.

Parameters *binary-spec* — Specifies the NAS-Port attribute

Values

binary-spec	<bit-specification> <binary-spec>
bit-specification	0 | 1 | <bit-origin>
bit-origin	*<number-of-bits><origin>
number-of-bits	1 to 32
origin	o | i | s | m | p
	outer VLAN ID
	i inner VLAN ID
	s slot number
	m MDA number
	p port number or lag-id

Output The following is an example of NAS port information.

Sample

```
*12o*12i00*2s*2m*2p => oooo oooo oooo iiii iiii iiii 00ss mmpp
If outer vlan = 0 & inner vlan = 1 & slot = 3 & mda = 1 & port = 1
=>  0000 0000 0000 0000 0000 0001 0011 0101 => nas-port = 309
```

nat-port-range

Syntax [**no**] **nat-port-range**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the of nat-port-range attribute and only applies to the 7750 SR.

The **no** form of this command reverts to the default.

remote-id

Syntax [**no**] **remote-id**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the agent-remote-id for RADIUS.

The **no** form of this command reverts to the default.

sap-session-index

Syntax [**no**] **sap-session-index**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute
config>subscr-mgmt>auth-plcy>include-radius-attribute

Description This command enables the generation of the per-SAP unique session index.

The **no** form of this command excludes **sap-session-index** attributes.

tunnel-server-attrs

Syntax [**no**] **tunnel-server-attrs**

Context config>subscr-mgmt>auth-policy>include-radius-attribute

Description This command includes tunnel-server attribute.

The **no** form of this command disables the generation of the tunnel-server-attrs RADIUS attribute.

wifi-num-attached-ues

Syntax [**no**] **wifi-num-attached-ues**

Context config>subscr-mgmt>acct-plcy>include
config>subscr-mgmt>auth-plcy>include

Description When enabled, this command indicates the number of UEs connected to the tunnel to which the radius message applies to. For session/host accounting this is the tunnel of the associated UE. For queue-instance accounting this attribute will not be included.

The **no** form of this command disables including the Alc-Num-Attached-UEs RADIUS attribute.

wifi-ssid-vlan

Syntax [**no**] **wifi-ssid-vlan**

Context config>subscr-mgmt>acct-plcy>include
config>subscr-mgmt>auth-plcy>include
config>aaa>isa-radius-plcy>auth-include-attributes
config>aaa>isa-radius-plcy>acct-include-attributes

Description This command enables including the per-SSID VLAN ID in a Alc-Wlan-SSID-VLAN.

The **no** form of this command disables including the per-SSID VLAN ID in Alc-Wlan-SSID-VLAN

sla-profile

Syntax [**no**] **sla-profile**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute
config>service>ies>sub-if>grp-if>sap>static-host
config>service>vprn>sub-if>grp-if>sap>static-host

Description This command specifies that SLA profile attributes should be included into RADIUS accounting messages.

The **no** form of this command reverts to the default.

spi-sharing-id

Syntax [**no**] **spi-sharing-id**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables RADIUS accounting messages to include the Alc-SPI-Sharing-Id RADIUS attribute. Together with the SLA profile name, this attribute provides details on the applicable SPI or queuing instance for this accounting session.

The **no** form of this command reverts to the default.

std-acct-attributes

Syntax [**no**] **std-acct-attributes**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables reporting of aggregated forwarded IPv4 and IPv6 octet, packet and gigaword counters using standard RADIUS attributes. This attribute is by default. It can be enabled simultaneously with detailed per queue or policer counters (**detailed-acct-attributes**).

The **no** form of this command reverts to the default.

sub-profile

Syntax [**no**] **sub-profile**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command specifies that subscriber profile attributes should be included into RADIUS accounting messages.

The **no** form of this command reverts to the default.

subscriber-id

Syntax [**no**] **subscriber-id**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command specifies that subscriber ID attributes should be included into RADIUS accounting messages.

The **no** form of this command reverts to the default.

tunnel-server-attrs

Syntax [**no**] **tunnel-server-attrs**

Context config>subscr-mgmt>auth-policy>include-radius-attribute
config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the generation of the tunnel-server RADIUS attribute and only applies to the 7750 SR.

The **no** form of this command reverts to the default.

uli

Syntax [**no**] **uli**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command, in case of ESM over GTP access, includes the ULI VSA in accounting. This VSA contains the last VSA as received via GTP. To generate a triggered interim update whenever ULI changes, the **config>subscr-mgmt>acct-plcy>triggered-updates>gtp-mobility** command is used.

The **no** version of this command disables inclusion of the ULI VSA.

Default no uli

user-name

Syntax [**no**] **user-name**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the inclusion of the user-name attribute.

The **no** form of this command disables the inclusion of the user-name attribute.

v6-aggregate-stats

Syntax [**no**] **v6-aggregate-stats**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables reporting of IPv6 aggregated forwarded octet and packet counters using RADIUS VSAs. Disabled by default. It requires **stat-mode v4-v6** for policers and queues for which the IPv6 aggregate forwarded packets should be counted.

The **no** form of this command reverts to the default.

wifi-rssi

Syntax [**no**] **wifi-rssi**

Context config>subscr-mgmt>acct-plcy>include-radius-attribute

Description This command enables the inclusion of the 802.11 Received Signal Strength Indication attribute.

The **no** form of this command reverts to the default.

password

Syntax **password** *password* [**hash** | **hash2**]
no password

Context config>subscr-mgmt>auth-policy

Description This command sets a password that is sent with **user-name** in every RADIUS authentication request sent to the RADIUS server upon receipt of DHCP discover or request messages. If no password is configured, no password AVP is sent.

The **no** form of this command reverts to the default value.

Parameters *password* — Specifies a text string containing the password. Allowed values are any string up to 64 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

hash — Specifies that the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2 — Specifies that the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

password

Syntax **password** *password* [**hash** | **hash2**]
no password

Context config>subscr-mgmt>diam-appl-plcy>nasreq

Description This command sets a password that is sent with **user-name** in every RADIUS authentication request sent to the RADIUS server upon receipt of DHCP discover or request messages. If no password is provided, an empty password is sent.

The **no** form of this command reverts to the default value.

Parameters *password* — Specifies a text string containing the password. Allowed values are any string up to 64 characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

hash — Specifies that the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2 — Specifies that the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

ppp-user-name

Syntax **ppp-user-name append** *domain-name*

ppp-user-name default-domain *domain-name*
ppp-user-name replace *domain-name*
ppp-user-name strip
no ppp-user-name

Context config>subscr-mgmt>auth-plcy

Description This command configures the password that is sent with the User-Name in Diameter NASREQ AA-Requests for IPoE hosts.

When no password is configured, an empty password is sent.

The **no** form of this command reverts to the default.

Parameters *domain-name* — A string specified by tmnxSubAuthPlcyPppDomain, preceded with a '@', is appended to the PAP/CHAP user name

domain-name — The same action is performed as with appendDomain, but only if the PAP/CHAP user name does not already contain a domain name

domain-name — All characters after a '@' delimiter are replaced with the string specified by tmnxSubAuthPlcyPppDomain

strip — Any “@” character and all subsequent characters are removed from the PAP/CHAP user name

pppoe-access-method

Syntax **pppoe-access-method** {**none** | **padi** | **pap-chap**}
no pppoe-access-method

Context config>subscr-mgmt>auth-plcy

Description This command indicates the authentication method used towards the RADIUS server in case the policy is used for PPPoE.

The **no** form of this command reverts to the default.

Parameters **none** — Indicates that the client is authenticated by the local user database defined under the group interface and not through RADIUS

padi — Indicates that the client is authenticated by RADIUS as soon as the PADI packet comes in (there is no PPP authentication done in the session in this case)

pap-chap — Indicates that the RADIUS authentication of the client is delayed until the authentication protocol phase in the PPP session (PAP or CHAP) and authentication is performed with the user name and PAP password / CHAP response supplied by the client

queue-instance-accounting

Syntax **queue-instance-accounting** [**interim-update**]
no queue-instance-accounting

Context config>subscr-mgmt>acct-plcy

Description This command enables per queue-instance-accounting. A stream of accounting messages (START/INTERIM-UPDATE/STOP) is generated per queuing instance. A queuing instance is equivalent to an sla-profile instance on non HSMDA based hardware and to subscriber on HSMDA based hardware. Accounting session id is generated per queuing instance and this accounting session id cannot be included in RADIUS Access-Request message. Queue instance counters represent volume based aggregation for all hosts sharing the queuing instance.

CoA and LI is supported based on the acct-session-id of the queuing instance.

The **no** form of this command reverts to the default.

Default queue-instance-accounting interim-update

Parameters **interim-update** — Specifies whether accounting messages are sent for the queue-instance. The queue-instance is the subscriber on High Scale MDA (HSMDA), or the SLA profile instance

radius-accounting-server

Syntax **radius-accounting-server**

Context config>subscr-mgmt>acct-plcy

Description This command enables the context to define RADIUS accounting server attributes under a given session accounting policy.

radius-accounting-server

Syntax **radius-authentication-server**

Context config>subscr-mgmt>acct-plcy

Description This command enables the context to define RADIUS accounting server attributes under a given session authentication policy.

access-algorithm

Syntax **access-algorithm** {**direct** | **round-robin**}

no access-algorithm

Context config>subscr-mgmt>auth-plcy-srvr
config>subscr-mgmt>acct-plcy>server

Description This command configures the algorithm used to access the list of configured RADIUS servers.

The **no** form of this command reverts to the default.

Parameters **direct** — Specifies that the first server is used as primary server for all requests, the second as secondary and so on

round-robin — Specifies that the first server is used as primary server for the first request, the second server as primary for the second request, and so on. If the router gets to the end of the list, it starts again with the first server.

hold-down-time

Syntax **hold-down-time** *seconds*
no hold-down-time

Context config>subscr-mgmt>auth-plcy>radius-auth-server

Description This command determines the interval during which no new communication attempts is made to a RADIUS server that is marked **down** to prevent immediately overloading the server when it is starting up. The only exception is when all servers in the authentication policy are marked **down**; in that case they will all be used again to prevent failures on new client connections.

The **no** form of this command reverts to the default.

Default hold-down 30

Parameters *seconds* — Specifies the hold time before re-using a RADIUS server that was down

Values 30 to 900

router

Syntax **router** *router-instance*
router service-name *service-name*
no router

Context config>subscr-mgmt>auth-plcy>radius-auth-server
config>subscr-mgmt>acct-plcy>server

Description This command specifies the virtual router instance applicable for the set of configured RADIUS servers. This value cannot be changed once a RADIUS server is configured for this policy. When the value is zero, both base and management router instances are matched.

The **no** form of this command reverts to the default.

Parameters *router-instance* — Specifies the virtual router instance

Values
router-name: Base, management
service-id: 1 to 2147483647

service-name — Specifies the service name up to 64 characters.

retry

Syntax **retry** *count*
no retry

Context config>subscr-mgmt>auth-plcy-srvr
config>subscr-mgmt>acct-plcy>server

Description This command configures the number of times the router attempts to contact the RADIUS server for authentication, if not successful the first time.

The **no** form of this command reverts to the default value.

Default retry 3

Parameters *count* — Specifies the retry count

Values 1 to 10

radius-server-policy

Syntax **radius-server-policy** *radius-server-policy-name*
no radius-server-policy

Context config>subscr-mgmt>auth-plcy
config>subscr-mgmt>acct-plcy

Description This command references an existing radius-server-policy (available under the config>aaa context) for use in subscriber management authentication and accounting.

When configured in an authentication-policy, following CLI commands are ignored in the policy to avoid conflicts:

- all commands in the radius-authentication-server context
- accept-authorization-change
- coa-script-policy
- accept-script-policy

- request-script-policy

When configured in a radius-accounting-policy, following CLI commands are ignored in the policy to avoid conflicts:

- all commands in the radius-accounting-server context
- acct-request-script-policy

The **no** form of this command removes the radius-server-policy reference from the configuration

Parameters *radius-server-policy-name* — Specifies the RADIUS server policy

server

Syntax **server** *server-index* **address** *ip-address* **secret** *key* [**hash** | **hash2**] [**port** *port-num*] [**coa-only**] [**pending-requests-limit** *limit*]
no server *index*

Context config>subscr-mgmt>auth-policy>radius-auth-server
config>subscr-mgmt>acct-plcy>server

Description This command adds a RADIUS server and configures the RADIUS server IP address, index, and key values.

Up to sixteen RADIUS servers can be configured at any one time in a RADIUS authentication policy. Only five can be used for authentication, all other servers should be configured as coa-only servers. In a RADIUS accounting policy, up to five RADIUS servers can be configured. RADIUS servers are accessed in order from lowest to highest index for authentication or accounting requests until a response from a server is received. A higher indexed server is only queried if no response is received from a lower indexed server (which implies that the server is not available). If a response from a server is received, no other RADIUS servers are queried.

The **no** form of this command removes the server index from the configuration.

Parameters *server-index* — Specifies the index for the RADIUS server. The index determines the sequence in which the servers are queried for authentication requests. Servers are queried in order from lowest to highest index.

Values 1 to 16 (a maximum of 5 authentication servers)

ip-address — Specifies the IP address of the RADIUS server. Two RADIUS servers cannot have the same IP address. An error message is generated if the server address is a duplicate.

key — Specifies the secret key to access the RADIUS server. This secret key must match the password on the RADIUS server.

Values secret-key: Up to 20 characters.
hash-key: Up to 33 characters.

hash2-ke: Up to 55 characters.

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

port-num — Specifies the UDP port number on which to contact the RADIUS server for authentication

Values 1 to 65535

coa-only — Specifies Change-of-Authorization Messages only. Servers that are marked with the coa-only flag will not be used for authentication, but they is able to accept RADIUS CoA messages, independent of the accept-authorization-change setting in the authentication policy.

For authentication purposes, the maximum number of servers is 5. All other servers may only be used as coa-only servers.

limit — Specifies the maximum number of outstanding RADIUS authentication requests for this authentication server

Values 1 to 4096

Default The default value when not configured is 4096

hold-down-time

Syntax [**no**] **hold-down-time** [**sec** *seconds*] [**min** *minutes*] [**hrs hours**] [**days** *days*]

Context config>aaa>radius-srv-plcy>servers

Description This command determines the interval during which no new communication attempts are made to a RADIUS server that is marked down to prevent immediately overloading the server when it is starting up. The only exception is when all servers in the authentication policy are marked down; in that case, they will all be used again to prevent failures on new client connections.

The **no** form of this command reverts to the default.

Default hold-down-time sec 30

Parameters *days* — Specifies the hold time in days before re-using a RADIUS server that was down

Values 1 to 1

hours — Specifies the hold time in hours before re-using a RADIUS server that was down

Values 1 to 23

minutes — Specifies the hold time in minutes before re-using a RADIUS server that was down

Values 1 to 59

seconds — Specifies the hold time in seconds before re-using a RADIUS server that was down

Values 1 to 59

down-timeout

Syntax [**no**] **down-timeout**

Context config>aaa>radius-srv-plcy>servers>health-check

Description This command determines the interval to wait for a RADIUS reply message from the RADIUS server before a RADIUS server is declared out-of-service. By default, the value of the down-timeout is the number of retries multiplied by the timeout interval. Each host will use the configured timeout and retry value under the AAA RADIUS server policy.

timeout refers to the waiting period before the next retry attempt

retry refers the number of times the host will attempt to contact the RADIUS server.

If a RADIUS server is declared out-of-service, the host pending retry attempts will move on to the next RADIUS server.

The **no** form of this command reverts to the default.

Parameters *minutes* — Specifies the timer to wait, in minutes, before declaring the RADIUS server that is down

Values 0 to 59

seconds — Specifies the timer to wait, in seconds, before declaring the RADIUS server that is down

Values 1 to 5

source-address

Syntax **source-address** *ip-address*
no source-address

Context config>subscr-mgmt>auth-plcy-srvr
config>subscr-mgmt>acct-plcy>server

Description This command configures the source address of the RADIUS packet.

The system IP address must be configured in order for the RADIUS client to work. See Configuring a System Interface in the Router Configuration Guide.

Note: The system IP address must only be configured if the source-address is not specified. When the **no source-address** command is executed, the source address is determined at the moment the request is sent. This address is also used in the nas-ip-address attribute: over there it is set to the system IP address if **no source-address** was given.

The **no** form of this command reverts to the default value.

Default System IP address

Parameters *ip-address* — Specifies the IP prefix for the IP match criterion in dotted decimal notation

Values 0.0.0.0 to 255.255.255.255

timeout

Syntax **timeout** *seconds*
no timeout

Context config>subscr-mgmt>auth-plcy-srvr
config>subscr-mgmt>acct-plcy>server

Description This command configures the number of seconds the router waits for a response from a RADIUS server.

The **no** form of this command reverts to the default value.

Default timeout 3

Parameters *seconds* — Specifies the time, in seconds, that the router waits for a response from a RADIUS server, expressed as a decimal integer

Values 1 to 90

session-accounting

Syntax **session-accounting** [**interim-update**] [**host-update**]
no session-accounting

Context config>subscr-mgmt>acct-plcy

Description This command enables per session accounting mode. In per session accounting mode, the acct-session-id is generated per session. This acct-session-id is uniformly included in all accounting messages (START/INTERIM-UPDATE/STOP) and it can be included in RADIUS Access-Request message.

This accounting mode of operation can be used only in PPPoE environment with dual-stack host in which case both hosts (IPv4 and IPv6) are considered part of the same session. In addition to regular interim-updates, *triggered* interim-updates are sent by a host joining or leaving the session.

When an IPv4/v6 address is allocated, or released from a dual-stack host, a triggered interim-update message is immediately sent. This triggered interim-update message reflects the change in the IP address. The triggered interim-update has no effect on the interval at which the regular interim updates are scheduled.

Accounting counters are based on the queue counters and as such are aggregated for all host sharing the queues within an sla-profile instance (non HSMDA) or a subscriber (HSMDA).

CoA and LI is supported based on the acct-session-id of the session.

The **no** form of this command reverts to the default.

Parameters **interim-update** — Specifies to enable interim updates. Without this keyword only START and STOP accounting messages are generated when the session is established or terminated. This is equivalent to a time-based accounting where only the duration of the session is required.

host-update — Indicates that host updates messages are sent. INTERIM-UPDATE messages can be generated (volume based accounting) by selecting this keyword

session-id-format

Syntax **session-id-format** {**description** | **number**}
no session-id-format

Context config>subscr-mgmt>acct-plcy

Description This command specifies the format for the acct-session-id attribute used in RADIUS accounting requests.

The **no** form of this command reverts to the default.

Default session-id-format description

Parameters **description** — Specifies to use a string containing following information: <subscriber>@<sap-id>@<SLA-profile>_<creation-time>

number — Specifies to use a unique number generated by the OS to identify a given session

triggered-updates

Syntax **triggered-updates**

Context config>subscr-mgmt>acct-plcy

Description This command enters the **triggered-updates** context. This node contains settings related to non-periodic accounting updates that are triggered by specific events.

gtp-mobility

Syntax [**no**] **gtp-mobility**

Context config>subscr-mgmt>acct-plcy>triggered-updates

Description This command, when enabled causes an interim update to be sent whenever a mobility event for ESM over GTP is detected. Mobility events include X2 mobility, service re-activation after idling, notification changes and any detected change of ULI or RAT-Type.

The **no** version of this command disables generating mobility related events.

Default no gtp-mobility

update-interval

Syntax **update-interval** *minutes*
no update-interval

Context config>subscr-mgmt>acct-plcy

Description This command specifies the interval at which accounting data of subscriber hosts is updated in a RADIUS Accounting Interim-Update message. Requires interim-update to be enabled when specifying the accounting mode in the radius accounting policy.

A RADIUS specified interim interval (attribute [85] Acct-Interim-Interval) overrides the CLI configured value.

The **no** form of this command reverts to the default.

Parameters *minutes* — Specifies the interval, in minutes, at which accounting data of subscriber hosts is updated

Values 5 to 259200

update-interval-jitter

Syntax **update-interval-jitter absolute** *seconds*

no update-interval-jitter

Context config>subscr-mgmt>acct-plcy

Description This command specifies the absolute maximum random delay introduced on the update interval between two accounting interim update messages. The effective maximum random delay value is the minimum of the configured absolute jitter value and 10% of the configured update-interval.

A value of zero will send the accounting interim update message without introducing an additional random delay.

The **no** form of this command sets the default to 10% of the configured update-interval.

Parameters *seconds* — Specifies the absolute maximum jitter value in seconds

Values 0 to 36000

re-authentication

Syntax [**no**] **re-authentication**

Context config>subscr-mgmt>auth-policy

Description This command enables authentication process at every DHCP address lease renewal s only if RADIUS did not reply any special attributes (for example, authentication only, no authorization).

The **no** form of this command reverts to the default value.

request-script-policy

Syntax **request-script-policy** *policy-name*
no request-script-policy

Context config>subscr-mgmt>auth-policy

Description This command specifies the RADIUS script policy used to change the RADIUS attributes of the outgoing Access-Request messages.

The **no** form of this command reverts to the default.

Parameters *policy-name* — Specifies a Python script policy to modify Access-Request messages

send-acct-stop-on-fail

Syntax **send-acct-stop-on-fail** {[**on-request-failure**] [**on-reject**] [**on-accept-failure**]}

no send-acct-stop-on-fail

Context config>subscr-mgmt>auth-policy

Description This command activates the reporting of RADIUS authentication failures of a PPPoE session to a RADIUS accounting server with an Accounting Stop message.

Three failure categories can be enabled separately:

- **on-request-failure**: All failure conditions between the sending of an Access-Request and the reception of an Access-Accept or Access-Reject.
- **on-reject**:
- **on-accept-failure**: All failure conditions that appear after receiving an Access-Accept and before successful instantiation of the host or session.

The RADIUS accounting policy to be used for sending the Accounting Stop messages must be obtained prior to RADIUS authentication via local user database pre-authentication.

The **no** form of this command reverts to the default.

Parameters **on-request-failure** — Specifies that an accounting stop message is sent when a RADIUS Access-Request message could not be sent (for example, there is no server configured, or timeout)

on-reject — Specifies that an accounting stop message is sent when an Access-Reject is received

on-accept-failure — Specifies that an accounting stop message is sent a failure occurred after the reception of a RADIUS Access-Accept message (such as a duplicate IP address)

user-name-format

Syntax **user-name-format** *format* [**mac-format** *mac-format*]
user-name-format *format* **append** [*domain-name*] [**mac-format** *mac-format*]
user-name-format *format* **append** *domain-name*
user-name-format *format* **default-domain** *domain-name* [**mac-format** *mac-format*]
user-name-format *format* **replace** *domain-name* [**mac-format** *mac-format*]
user-name-format *format* **strip** [**mac-format** *mac-format*]
no user-name-format

Context config>subscr-mgmt>auth-policy

Description This command defines the format of the “user-name” field in the session authentication request sent to the RADIUS server.

The **no** form of this command switches to the default format, **mac**.

Default By default, the MAC source address of the DHCPv4 DISCOVER or DCHPv6 SOLICIT message is used in the user-name field.

Parameters *format* — Specifies the user name format in RADIUS message

Values ascii-converted-circuit-id, ascii-converted-tuple, circuit-id, dhcp-client-vendor-opts, mac, mac-giaddr, tuple

ascii-converted-circuit-id — Identical to circuit-id, but the user name is sent to the RADIUS server as a string of hex digits, for use if there is binary data in the circuit-id

ascii-converted-tuple — Identical to tuple, but the circuit-id part of the user name is sent to the RADIUS server as a string of hex digits, for use if there is binary data in the circuit-id

circuit-id — If the system serves as a DHCP relay server which inserts option 82 info, the user name is formatted as defined under DHCP information option.
If the system is not a DHCP relay server, the circuit-id is taken from option 82 in the received DHCP message. If no circuit-id can be found, the DHCP-msg is rejected.

dhcp-client-vendor-opts — IPoEv4 host (IPoE session enabled or disabled on group-interface) — The RADIUS user-name is a concatenation of the DHCPv4 Client Identifier Option 61, an "@" delimiter, and the DHCPv4 Vendor Class Identifier Option 60. Non-printing characters in the DHCP option values are converted as described below.

IPoEv6 host (IPoE session enabled on group-interface) — The RADIUS user-name is a concatenation of the identifier field of a type 2 DUID in the DHCPv6 Client Identifier Option 1, the "@" delimiter, and the opaque data field of the first vendor class data in the DHCPv6 Vendor Class Option 16. Non-printing characters in the DHCP option values are converted as described below.

IPoEv6 host (IPoE session disabled on group-interface) — The MAC source address of the DHCPv6 SOLICIT message.

In the absence of a DHCPv4 Client Identifier Option 61 or a DHCPv6 Client Identifier Option 1 containing a DUID type 2, the DHCP client MAC address is used.

In the absence of a DHCPv4 Vendor Class Identifier Option 60 or a DHCPv6 Vendor Class Option 16, the "@" delimited is omitted and nothing is appended.

Non-printing characters, that is, characters outside the ASCII range hex 21 through hex 7E, are converted into an underscore (hex 5F) character.

mac — The MAC source address of the DHCPv4 DISCOVER or DHCPv6 SOLICIT message is used in the user-name field. The format of the MAC address string used as the user name in the RADIUS authentication requests uses lowercase hex digits, and ":" as the inter-digit separator, for example, 00:11:22:aa:bb:cc is valid but 00-11-22-AA-BB-CC will return an error. The RADIUS server must be configured accordingly, otherwise the authentication request will fail.

mac-giaddr — A concatenation of the MAC address and the Relay Agent IP address (giaddr)

tuple — Specifies that the concatenation of MAC source address and circuit-ID are used in the user-name field

mac-format — Specifies how a MAC address is represented when contacting a RADIUS server. This is only used while the value of is equal to the DHCP client vendor options and if the MAC address is used by default of the DHCP client vendor options.

Examples: ab: 00:0c:f1:99:85:b8 7xxx style

XY-	00-0C-F1-99-85-B8 IEEE canonical style
mmmm.	0002.03aa.abff Cisco style

append — Specifies the data type which is an enumerated integer that indicates what needs to be appended to the user-name sent to the RADIUS server

Values 1 — nothing
2 — domain name

domain — Specifies to user the domain name. In some instances it is desired to add a domain only to usernames which have omitted the domain (@domain). In these instances a default-domain can be appended to usernames which lack a @domain.

append — Adds a "@" delimiter and the specified string after the PAP/CHAP username. No allowance is made for the presence of an existing domain or @ delimited.

replace — Replaces the character-string after the "@" delimiter with the string specified

strip — Removes all characters after and including the "@" delimiter.

For example:

```
Command: append
String:     domainA-1.com
PAP/CHAP User:    someuser
Resulting User:    someuser@domainA-1.com

Command: append
String:     domainA-1.com
PAP/CHAP User:    someuser@existing-domain.net
Resulting User:    someuser@existing-domain.net@domainA-1.com

Command: strip
String:
PAP/CHAP User:    someuser@existing-domain.net
Resulting User:    someuser

Command: replace
String:     domainA-1.com
PAP/CHAP User:    someuser@existing-domain.net
Resulting User:    someuser@domainA-1.com

Command: default-domain
String:    domainA-1.com
PAP/CHAP User:    someuser@existing-domain.net
Resulting User:    someuser@existing-domain.net

Command: default-domain
String: domainA-1.com
PAP/CHAP User:    someuser
Resulting User:    someuser@domainA-1.com
```

user-name-format

Syntax **user-name-format** *format*

no user-name-format

Context config>subscr-mgmt>diam-appl-plcy>nasreq

Description This command defines the format of the User-Name AVP value in Diameter NASREQ AA-Requests for IPoE hosts.

The **no** form of this command reverts to the default.

Parameters *format* — Specifies the format of the User-Name AVP value

Values

mac — Specifies to use the MAC source address of the DHCPv4 DISCOVER or DHCPv6 SOLICIT message in the user-name field. The format of the MAC address string is defined with the **mac-forma**t CLI command

circuit-id — Specifies to use the circuit ID to identify the user toward the server. If the system serves as a DHCP relay server which inserts option 82 info, the user name is formatted as defined under DHCP information option. If the system is not a DHCP relay server, the circuit-id is taken from option 82 in the received DHCP message. If no circuit-id can be found, the DHCP-msg is rejected.

tuple — Specifies to use a concatenation of MAC source address and circuit-ID

ascii-converted-circuit-id — Identical to circuit-id, but the user name is a string of hex digits, for use if there is binary data in the circuit-id

ascii-converted-tuple — Specifies a string of hex digits, for use if there is binary data in the circuit-id

dhcp-client-vendor-opts — Specifies to use a concatenation of the DHCPv4 Client Identifier Option 61, the “@” delimiter, and the DHCPv4 Vendor Class Identifier Option 60. Non-printing characters in the DHCP option values are converted as described below.

In the absence of a DHCPv4 Client Identifier Option 61, the DHCP client MAC address is used.

In the absence of a DHCPv4 Vendor Class Identifier Option 60, the “@” delimiter is omitted and nothing is appended.

Non-printing characters, that is, characters outside the ASCII range hex 21 through hex 7E, are converted into an underscore (hex 5F) character.

mac-giaddr — Specifies to use a concatenation of MAC source address and DHCP GI address

nas-port-id — Specifies to use a value of the nas-port-id with format defined in the include-avp section

user-name-operation

Syntax **user-name-operation** *operation* [**domain** *domain-name*]
no user-name-operation

Context config>subscr-mgmt>diam-appl-plcy>nasreq

Description This command enables domain name manipulation of the user name, such as append, strip, replace or add as default.

For IPoE, this command only applies when **user-name-format** is configured to **dhcp-client-vendor-opts**.

The **no** form of this command reverts to the default.

Parameters *operation* — Specifies the user name manipulations with respect to domain name values

Values **append-domain** – appends an "@" delimiter with the specified domain-name at the end of the user-name, independent if a domain name was already present

strip-domain – removes all characters after and including the "@" delimiter

default-domain – adds an "@" delimiter and the specified domain name to user-names that have no domain name present

replace-domain – replaces the characters after the "@" delimiter with the specified domain-name

domain-name — Specifies the domain name string to be used in the specified operation, up to 128 characters

9.25.2.1.4 BGP Peering Policy Commands

Note: The commands described in this section apply only to the 7750 SR.

bgp-peering-policy

Syntax **bgp-peering-policy** *policy-name* [**create**]
no bgp-peering-policy *policy-name*

Context config>subscr-mgmt

Description This command configures the name of the BGP peering policy.

The **no** form of this command removes the policy name from the configuration.

Parameters *policy-name* — Specifies the BGP peer policy name, up to 32 characters

create — Keyword used to create the peering policy. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

advertise-inactive

Syntax [**no**] **advertise-inactive**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command enables the advertising of inactive BGP routers to other BGP peers.

By default, BGP only advertises BGP routes to other BGP peers if a given BGP route is chosen by the route table manager as the most preferred route within the system and is active in the forwarding plane. This command allows system administrators to advertise a BGP route even though it is not the most preferred route within the system for a given destination.

The **no** form of this command disables the advertising.

aggregator-id-zero

Syntax [**no**] **aggregator-id-zero**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command is used to set the router ID in the BGP aggregator path attribute to zero when BGP aggregates routes. This prevents different routers within an AS from creating aggregate routes that contain different AS paths.

When BGP is aggregating routes, it adds the aggregator path attribute to the BGP update messages. By default, BGP adds the AS number and router ID to the aggregator path attribute.

When this command is enabled, BGP adds the router ID to the aggregator path attribute. The **no** form of this command used at the global level reverts to default where BGP adds the AS number and router ID to the aggregator path attribute.

as-override

Syntax [**no**] **as-override**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command replaces all instances of the peer's AS number with the local AS number in a BGP route's AS_PATH.

This command breaks BGP's loop detection mechanism. It should be used carefully.

The **no** form of this command reverts to the default.

auth-keychain

Syntax **auth-keychain** *name*
no auth-keychain

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the BGP authentication key for all peers.

The keychain allows the rollover of authentication keys during the lifetime of a session.

The **no** form of this command reverts to the default.

Parameters *name* — Specifies the name of an existing keychain, up to 32 characters, to use for the specified TCP session or sessions

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the BGP authentication key.

Authentication is performed between neighboring routers before setting up the BGP session by verifying the password. Authentication is performed using the MD-5 message-based digest. The authentication key can be any combination of letters or numbers from 1 to 16.

The **no** form of this command removes the authentication password from the configuration and effectively disables authentication.

Parameters *authentication-key* — Specifies an authentication key. The key can be any combination of ASCII characters, up to 255 characters (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

hash-key — Species the hash key. The key can be any combination of ASCII characters up to 342 characters (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

cluster

Syntax **cluster** *cluster-id*
no cluster

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the cluster ID for a route reflector server.

Route reflectors are used to reduce the number of IBGP sessions required within an AS. Normally, all BGP speakers within an AS must have a BGP peering with every other BGP speaker in an AS. A route reflector and its clients form a cluster. Peers that are not part of the cluster are considered to be non-clients.

When a route reflector receives a route, first it must select the best path from all the paths received. If the route was received from a non-client peer, then the route reflector sends the route to all clients in the cluster. If the route came from a client peer, the route reflector sends the route to all non-client peers and to all client peers except the originator.

For redundancy, a cluster can have multiple route reflectors.

Confederations can also be used to remove the full IBGP mesh requirement within an AS.

The **no** form of this command deletes the cluster ID and effectively disables the Route Reflection for the given group.

Parameters *cluster-id* — Specifies the route reflector cluster ID is expressed in dot decimal notation

Values Any 32 bit number in dot decimal notation. (0.0.0.1 to 255.255.255.255)

connect-retry

Syntax **connect-retry** *seconds*
no connect-retry

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the BGP connect retry timer value in seconds.

When this timer expires, BGP tries to reconnect to the configured peer.

The **no** form of this command used at the global level reverts to the default value.

Default connect-retry 120

Parameters *seconds* — The BGP Connect Retry timer value in seconds, expressed as a decimal integer

Values 1 to 65535

damping

Syntax [**no**] **damping**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command enables BGP route damping for learned routes which are defined within the route policy. Use damping to reduce the number of update messages sent between BGP peers and reduce the load on peers without affecting the route convergence time for stable routes. Damping parameters are set via route the policy definition.

The **no** form of this command used at the global level disables route damping.

When damping is enabled and the route policy does not specify a damping profile, the default damping profile is used. This profile is always present and consists of the following parameters:

Half-life:	15 minutes
Max-suppress:	60 minutes
Suppress-threshold:	3000
Reuse-threshold	750

disable-4byte-asn

Syntax [**no**] **disable-4byte-asn**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command disables the use of 4-byte ASNs. It can be configured at all 3 level of the hierarchy so it can be specified down to the per peer basis.

If this command is enabled 4-byte ASN support should not be negotiated with the associated remote peer.

The **no** form of this command resets the behavior to the default which is to enable the use of 4-byte ASN.

disable-client-reflect

Syntax [**no**] **disable-client-reflect**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command disables the reflection of routes by the route reflector to the group or neighbor. This only disables the reflection of routes from other client peers. Routes learned from non-client peers are still reflected to all clients.

The **no** form re-enables client reflection of routes to all client peers.

disable-communities

Syntax **disable-communities** [**standard**] [**extended**]
no disable-communities

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures BGP to disable sending communities.

The **no** form of this command reverts to the default.

Parameters **standard** — Specifies standard communities that existed before VPRNs or 2547

extended — Specifies BGP communities used were expanded after the concept of 2547 was introduced, to include handling the VRF target

disable-fast-external-failover

Syntax [**no**] **disable-fast-external-failover**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures BGP fast external failover.

The **no** form of this command reverts to the default.

export

Syntax **export** *policy* [*policy*]
no export

Context config>subscr-mgmt>bgp-prng-plcy

Description This command specifies the export policies to be used to control routes advertised to BGP neighbors.

When multiple policy names are specified, the policies are evaluated in the order they are specified. A maximum of five (5) policy names can be configured. The first policy that matches is applied.

Note: If a non-existent route policy is applied to a VPRN instance, the CLI generates a warning message. This message is only generated at an interactive CLI session and the route policy association is made. No warning message is generated when a non-existent route policy is applied to a VPRN instance in a configuration file or when SNMP is used.

The **no** form of this command removes all route policy names from the export list.

Default no export — BGP advertises routes from other BGP routes but does not advertise any routes from other protocols unless directed by an export policy

Parameters *policy* — Specifies a route policy statement name, up to 32 characters

hold-time

Syntax **hold-time** *seconds*
no hold-time

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the BGP hold time, expressed in seconds.

The BGP hold time specifies the maximum time BGP waits between successive messages (either keepalive or update) from its peer, before closing the connection.

Even though the router OS implementation allows setting the keepalive time separately, the configured keepalive timer is overridden by the hold-time value under the following circumstances:

If the specified hold-time is less than the configured keepalive time, then the operational keepalive time is set to a third of the hold-time; the configured keepalive time is not changed.

If the hold-time is set to zero, then the operational value of the keepalive time is set to zero; the configured keepalive time is not changed. This means that the connection with the peer is up permanently and no keepalive packets are sent to the peer.

The **no** form of this command reverts to the default.

Parameters *seconds* — Specifies the hold-time, in seconds, expressed as a decimal integer. A value of 0 indicates the connection to the peer is up permanently.

Values 0, 3 to 65535

import

Syntax **import** *policy* [*policy*]
no import

Context config>subscr-mgmt>bgp-prng-plcy

Description This command specifies the import policies to be used to control routes advertised to BGP neighbors. Route policies are configured in the **config>router>policy-options** context. When multiple policy names are specified, the policies are evaluated in the order they are specified. The first policy that matches is applied.

The **no** form of this command removes all route policy names from the import list.

Default no import — BGP accepts all routes from configured BGP neighbors. Import policies can be used to limit or modify the routes accepted and their corresponding parameters and metrics.

Parameters *policy* — Specifies route policy statement name up to 32 characters. Up to five policies can be specified.

keepalive

Syntax **keepalive** *seconds*
no keepalive

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the BGP keepalive timer. A keepalive message is sent every time this timer expires.

The keepalive value is generally one-third of the hold-time interval. Even though the OS implementation allows the keepalive value and the hold-time interval to be independently set, under the following circumstances, the configured keepalive value is overridden by the hold-time value:

If the specified keepalive value is greater than the configured hold-time, then the specified value is ignored, and the keepalive is set to one third of the current hold-time value.

If the specified hold-time interval is less than the configured keepalive value, then the keepalive value is reset to one third of the specified hold-time interval.

If the hold-time interval is set to zero, then the configured value of the keepalive value is ignored. This means that the connection with the peer is up permanently and no keepalive packets are sent to the peer.

The **no** form of this command used at the global level reverts to the default value.

Default keepalive 30

Parameters *seconds* — Specifies the keepalive timer in seconds, expressed as a decimal integer

Values 0 to 21845

local-address

Syntax **local-address** *ip-address*
no local-address

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the local IP address used by the group or neighbor when communicating with BGP peers.

Outgoing connections use the **local-address** as the source of the TCP connection when initiating connections with a peer.

When a local address is not specified, the 7750 SR OS uses the system IP address when communicating with IBGP peers and uses the interface address for directly connected EBGP peers. This command is used at the neighbor level to revert to the value defined under the group level.

The **no** form of this command removes the configured local-address for BGP.

The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Parameters *ip-address* — Specifies the IPv4 or IPv6 address of the local address

For IPv4, the local address is expressed in dotted decimal notation. Allowed values are a valid routable IP address on the router, either an interface or system IP address.

For IPv6, the local address is expressed in semi-colon hexadecimal notation. Allowed values is an interface or a system IP address.

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D

local-as

Syntax **local-as** *as-number* [**private**]
no local-as

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures a BGP virtual autonomous system (AS) number.

In addition to the AS number configured for BGP in the config>router>autonomous-system context, a virtual (local) AS number is configured. The virtual AS number is added to the as-path message before the router's AS number makes the virtual AS the second AS in the as-path.

This configuration parameter can be set at three levels: global level (applies to all peers), group level (applies to all peers in peer-group) or neighbor level (only applies to specified peer). Thus, by specifying this at each neighbor level, it is possible to have a separate as-number per EBGP session.

When a command is entered multiple times for the same AS, the last command entered is used in the configuration. The **private** attribute can be added or removed dynamically by reissuing the command.

Changing the local AS at the global level in an active BGP instance causes the BGP instance to restart with the new local AS number. Changing the local AS at the global level in an active BGP instance causes BGP to re-establish the peer relationships with all peers in the group with the new local AS number. Changing the local AS at the neighbor level in an active BGP instance causes BGP to re-establish the peer relationship with the new local AS number.

This is an optional command and can be used in the following circumstance:

Provider router P is moved from AS1 to AS2. The customer router that is connected to P, however, is configured to belong to AS1. To avoid reconfiguring the customer router, the **local-as** value on router P can be set to AS1. Thus, router P adds AS1 to the as-path message for routes it advertises to the customer router.

The **no** form of this command used at the global level will remove any virtual AS number configured.
The **no** form of this command used at the group level reverts to the value defined at the global level.

The **no** form of this command used at the neighbor level reverts to the value defined at the group level.

Parameters *as-number* — The virtual autonomous system number, expressed as a decimal integer

Values 1 to 65535

private — Specifies the local-as is hidden in paths learned from the peering

local-preference

Syntax **local-preference** *local-preference*
no local-preference

Context config>subscr-mgmt>bgp-prng-plcy

Description This command enables setting the BGP local-preference attribute in incoming routes if not specified and configures the default value for the attribute. This value is used if the BGP route arrives from a BGP peer without the **local-preference** integer set.

The specified value can be overridden by any value set via a route policy.

The **no** form of this command at the global level specifies that incoming routes with local-preference set are not overridden and routes arriving without local-preference set are interpreted as if the route had local-preference value of 100.

Parameters *local-preference* — The local preference value to be used as the override value, expressed as a decimal integer

Values 0 to 4294967295

loop-detect

Syntax **loop-detect** {**drop-peer** | **discard-route** | **ignore-loop** | **off**}
no loop-detect

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures how the BGP peer session handles loop detection in the AS path.

Note: Dynamic configuration changes of **loop-detect** are not recognized.

The **no** form of this command used at the global level reverts to default, which is **loop-detect ignore-loop**.

Parameters **drop-peer** — Sends a notification to the remote peer and drops the session

discard-route — Discards routes received with loops in the AS path

ignore-loop — Ignores routes with loops in the AS path but maintains peering

off — Disables loop detection

med-out

Syntax **med-out** {*number* | **igp-cost**}
no med-out

Context config>subscr-mgmt>bgp-prng-plcy

Description This command enables advertising the Multi-Exit Discriminator (MED) and assigns the value used for the path attribute for the MED advertised to BGP peers if the MED is not already set.

The specified value can be overridden by any value set with a route policy.

The **no** form of this command used at the global level reverts to default where the MED is not advertised.

Parameters *number* — Specifies the MED path attribute value, expressed as a decimal integer

Values 0 to 4294967295

igp-cost — Specifies that the MED is set to the IGP cost of the given IP prefix

min-route-advertisement

Syntax **min-route-advertisement** *seconds*
no min-route-advertisement

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the minimum interval, in seconds, at which a prefix can be advertised to a peer.

The **no** form of this command reverts to default values.

Default min-route-advertisement 30

Parameters *seconds* — Specifies the minimum route advertising interval, in seconds, expressed as a decimal integer

Values 1 to 255

multihop

Syntax **multihop** *ttl-value*
no multihop

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the time to live (TTL) value entered in the IP header of packets sent to an EBGP peer multiple hops away.

This parameter is meaningful only when configuring EBGP peers. It is ignored if set for an IBGP peer.

The **no** form of this command is used to convey to the BGP instance that the EBGP peers are directly connected.

The **no** form of this command reverts to default values.

Default multihop 1 (EBGP peers are directly connected)

multihop 64 (IBGP)

Parameters *ttl-value* — Specifies the TTL value, expressed as a decimal integer

Values 1 to 255

next-hop-self

Syntax [**no**] **next-hop-self**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the neighbor to always set the NEXTHOP path attribute to its own physical interface when advertising to a peer.

The **no** form of this command disables the command.

passive

Syntax [**no**] **passive**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command enables the passive mode for the BGP neighbors.

The **no** form of this command disables the passive mode.

peer-as

Syntax **peer-as** *as-number*
no peer-as

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the autonomous system number for the remote peer. The peer AS number must be configured for each configured peer.

The **no** form of this command removes the *as-number* from the configuration.

Parameters *as-number* — Specifies the AS number for the remote peer

Values 1 to 4294967295

preference

Syntax [**no**] **preference** *preference*

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the route preference for routes learned from the configured peer.

The lower the preference the higher the chance of the route being the active route. The OS assigns BGP routes highest default preference compared to routes that are direct, static or learned via MPLS or OSPF.

The **no** form of this command used at the global level reverts to default value.

Default preference 170

Parameters *preference* — The route preference, expressed as a decimal integer

Values 1 to 255

prefix-limit

Syntax **prefix-limit** *limit*
no prefix-limit

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the maximum number of routes BGP can learn from a peer.

When the number of routes reaches 90% of this limit, an SNMP trap is sent. When the limit is exceeded, the BGP peering is dropped and disabled.

The **no** form of this command removes the **prefix-limit**.

Parameters *limit* — The number of routes that can be learned from a peer, expressed as a decimal integer

Values 1 to 4294967295

remove-private

Syntax [**no**] **remove-private**

Context config>subscr-mgmt>bgp-prng-plcy

Description This command allows private AS numbers to be removed from the AS path before advertising them to BGP peers.

The OS software recognizes the set of AS numbers that are defined by IANA as private. These are AS numbers in the range 64512 through 65535, inclusive.

The **no** form of this command used at the global level reverts to default value.

Private AS numbers are included in the AS path attribute.

type

Syntax [**no**] **type** {**internal** | **external**}

Context config>subscr-mgmt>bgp-prng-plcy

Description This command designates the BGP peer as type internal or external.

The type of **internal** indicates the peer is an IBGP peer while the type of external indicates that the peer is an EBGP peer.

By default, the OS derives the type of neighbor based on the local AS specified. If the local AS specified is the same as the AS of the router, the peer is considered **internal**. If the local AS is different, then the peer is considered **external**.

The **no** form of this command used at the group level reverts to the default value.

Parameters **internal** — Configures the peer as internal

external — Configures the peer as external

ttl-security

Syntax **ttl-security** *min-ttl-value*
no ttl-security

Context config>subscr-mgmt>bgp-prng-plcy

Description This command configures the TTL security parameters for incoming packets.

The **no** form of this command reverts to the default.

Parameters *min-ttl-value* — Specifies the minimum TTL value for an incoming BGP packet

Values 1 to 255

9.25.2.1.5 Call Trace Commands

call-trace

Syntax **call-trace**

Context config

Description This command enables the context to configure parameters related to the call trace debugging tool.

location

Syntax **location** {**cf1** | **cf2**}

Context config>call-trace

Description This command specifies the compact flash (CF) configuration to store call trace files.

disable

Syntax [**no**] **disable**

Context config>call-trace>location

Description When configured, the specified compact flash will not be used by call-trace.

The **no** form of this command enables the compact flash for use by call-trace.

size-limit

Syntax **size-limit** *limit-value*
no size-limit

Context config>call-trace>location

Description This command limits the total size of call-trace files on the specified compact flash card.

The **no** form of this command removes the size restriction.

Default size-limit 1000

Parameters *limit-value* — Specifies the total size of call-trace files on the specified compact flash card, in Mbytes

Values 1 to 65536

max-files-number

Syntax **max-files-number** *number*

Context config>call-trace

Description This command configures the maximum number of files call trace can create.

Default max-files-number 200

Parameters *number* — Specifies the maximum number of all call trace log files stored on all compact flash cards together

Values 1 to 1024

primary-cf

Syntax **primary-cf** *cflash-id*

Context config>call-trace

Description This command specifies which compact-flash is used as the primary CF for call-trace operation.

Default primary-cf cf1

Parameters *cflash-id* — Specifies the compact flash card to be used as the primary local storage location to save the generated call trace log files

Values cf1, cf2

trace-profile

Syntax **trace-profile** *profile-name* [**create**]
no trace-profile *profile-name*

Context config>call-trace

Description This command creates a profile that can be applied to a specific trace job.

Parameters *profile-name* — Specifies the unique name of the call trace profile.

create — Keyword used to create the trace profile instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context

applications

Syntax **applications** [**connectivity-management**] [**radius-auth**] [**radius-acct**] [**python**] [**ludb**]

[**msap**]
applications all

Context config>call-trace>trace-profile

Description This command enables tracing of messages and events for the specified applications.

Default applications all

Parameters **all** — Enables tracing of all packets and events.

connectivity-management — Enables tracing for connectivity protocols, such as DHCP, ARP, and DHCPv6, and events related to connectivity management; for example, migrant or data-triggered host creation, idling, or session timeout

ludb — Enables tracing of local user database lookups.

msap — Enables tracing of MSAP creation events.

python — Enables tracing of Python script execution.

radius-acct — Enables tracing of messages and events related to RADIUS-based accounting.

radius-auth — Enables tracing of messages and events related to RADIUS authentication, including CoA and Disconnect.

debug-output

Syntax [**no**] **debug-output**

Context config>call-trace>trace-profile

Description This command enables output to the debug log. Traced messages are decoded as text and sent to debug logging. This allows for real-time debugging but is not recommended on live systems because of the decode overhead.

This command disables logging to the compact flash. The **debug-output** and **live-output** commands are mutually exclusive.

The **no** form of this command disables output to the debug log.

events

Syntax **events** {**none** | **public-only** | **all**}

Context config>call-trace-trace-profile

Description This command configures whether captured traces include events that occurred on the SR OS router, such as mobility and idle-timeout.

Default events none

Parameters **none** — Specifies that no events is traced

public-only — Specifies that only events that are readable by everyone is traced

all — Specifies that all events is traced, including events that are encrypted for use by customer support only. Encrypted events are not readable by end-users

live-output

Syntax **live-output** {*ip-address* | *fqdn*} [**port** *port*] [**router** {*router-instance* | **service-name** *service-name*}]
no live-output

Context config>call-trace>trace-profile

Description This command specifies a live output destination for this trace. When configured, captures will not be stored locally but sent (over UDP) to the server in the specified routing context. The destination can be specified as either an IP address or a DNS FQDN. The **live-output** and **debug-output** commands are mutually exclusive.

The **no** form of this command disables live output streaming.

Parameters *ip-address* — Specifies the IPv4 or IPv6 address of the server to stream to

fqdn — Specifies the FQDN that represents the server in DNS, 255 characters max

port — Specifies the UDP port on which the server is listening

Values 1 to 65535

Default 29770

router-instance — Specifies the router instance in which the live output is forwarded

service-name — Specifies the name of the Layer 3 service in which the live output is forwarded

size-limit

Syntax **size-limit** *limit-value*

Context config>call-trace>trace-profile

Description This command specifies a maximum of how big a trace may grow before it is stopped.

Default size-limit 10

Parameters *limit-value* — Specifies the maximum data volume generated by a single call trace job to the output in megabytes. After reaching the limit the call trace job for a given host is automatically terminated.

Values 1 to 1000

time-limit

Syntax **time-limit** *limit-value*

Context config>call-trace>trace-profile

Description This command specifies how long a trace may run before it is stopped.

Default time-limit 86400

Parameters *limit-value* — Specifies the maximum duration of a single call trace job in seconds. After reaching the limit the call trace job for a given host is automatically terminated.

Values 1 to 604800

9.25.2.1.6 Category Map Commands

category-map

Syntax **category-map** *category-map-name* [**create**]
no category-map *category-map-name*

Context config>subscr-mgmt
config>subscr-mgmt>sla-prof

Description This command specifies the category map name.

The **no** form of this command reverts to the default.

Parameters *category-map-name* — Specifies the category map name, up to 32 characters

create — Keyword used to create a category map instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

credit-control-policy

Syntax **credit-control-policy** *policy-name* [**create**]
credit-control-policy diameter *policy-name*
no credit-control-policy *policy-name*

Context config>subscr-mgmt

Description This command creates, configures or deletes a credit control policy.

The **no** form of this command reverts to the default.

Parameters *policy-name* — Specifies the policy name, up to 32 characters

create — Keyword used to create the credit control policy. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

credit-control-server

Syntax **credit-control-server radius**
credit-control-server diameter *policy-name*
no credit-control-server

Context config>subscr-mgmt>credit-control-policy

Description This command configures the credit control server to use. In case of RADIUS, the servers defined in the authentication policy are used. For Diameter, the peers defined in the specified Diameter policy are used.

The **no** form of this command reverts to the default.

Default credit-control-server radius

Parameters **radius** — Specifies to use the RADIUS authentication servers defined in the RADIUS authentication policy in the group interface to report credit usage and obtain new credit

diameter *policy-name* — Specifies to use the diameter peers specified in the diameter policy to report credit usage and obtain new credit

default-category-map

Syntax **default-category-map** *category-map-name*
no default-category-map

Context config>subscr-mgmt>credit-control-policy

Description This command configures the default category map.

The **no** form of this command reverts to the default.

Parameters *category-map-name* — Specifies the category map name, up to 32 characters

error-handling-action

Syntax **error-handling-action** {**continue** | **block**}
no error-handling-action

Context config>subscr-mgmt>credit-control-policy

Description This command configures the error handling action for the policy.

The **no** form of this command reverts to the default.

Default error-handling-action continue

Parameters **continue** — Specifies to continue when an error occurs in the CC determination

block — Specifies to block when an error occurs in the CC determination

out-of-credit-action

Syntax **out-of-credit-action** {**continue** | **disconnect-host** | **block-category** | **change-service-level**}
no out-of-credit-action

Context config>subscr-mgmt>credit-control-policy

Description This command configures the action to be performed when out of credit is reached.

The **no** form of this command reverts to the default.

Parameters *action* — Specifies the action to be taken when out of credit is out reached

Values **continue**, **disconnect-host**, **block-category**, **change-service-level**

activity-threshold

Syntax **activity-threshold** *kilobits-per-second*
no activity-threshold

Context config>subscr-mgmt>cat-map

Description This command configures the threshold that is applied to determine whether or not there is activity. This is only valid for credit-type = time (not volume).

The **no** form of this command reverts to the default.

Parameters *kilobits-per-second* — Specifies the activity threshold value in kilobits per second

Values 1 to 100000000

category

Syntax **category** *category-name* [**create**]
no category *category-name*

Context config>subscr-mgmt>cat-map

Description This command enables the context for use in RADIUS credit control, Diameter credit control (Gy), Diameter Gx Usage Monitoring, or Idle-Timeout.

Up to sixteen categories can be configured per category map. The internal category for Gx session level Usage Monitoring is included in this limit. The instantiation of the internal category is controlled with the **gx-session-level-usage** command.

Parameters *category-name* — Specifies the category name, up to 32 characters

create — Keyword used to create a category instance. The **create** keyword can be enabled or disabled in the **environment>create** context.

def-instance-sharing

Syntax **def-instance-sharing** {**per-sap** | **per-session**}

Context config>subscr-mgmt>sla-prof

Description This command configures the default SPI sharing method for IPoE or PPPoE sessions from the same subscriber, having the same SLA profile associated, that are active on the same SAP. SPI sharing can be per-SAP or a dedicated SPI per-session.

Default def-instance-sharing per-sap

Parameters **per-sap** — Specifies that IPoE or PPP sessions from the same subscriber, having the same SLA profile associated, and active on the same SAP share an SPI

per-session — Specifies that IPoE or PPP sessions from the same subscriber, having the same SLA profile associated, and active on the same SAP obtain a dedicated SPI per session

category-map

Syntax **category-map** *category-map-name*
no category-map

Context config>subscr-mgmt>sla-prof>idle-timeout

Description This command references the category-map to be used for the idle-timeout monitoring of subscriber hosts associated with this sla-profile. The **category-map** must already exist in the **config**>**subscr-mgmt** context.

The **no** form of this command reverts to the default.

Parameters *category-map-name* — Specifies the name of the category map, up to 32 characters, where the activity-threshold and the category is defined for idle-timeout monitoring of subscriber hosts

category

Syntax **category** *category-name* [**create**]
no category *category-name*

Context config>subscr-mgmt>sla-prof>cat-map

Description This command defines the category in the category-map to be used for the idle-timeout monitoring of subscriber hosts.

The **no** form of this command reverts to the default.

Parameters *category-name* — Specifies the name, up to 32 characters, of the category where the queues and policers are defined for idle-timeout monitoring of subscriber hosts

create — Keyword used to create a category instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

idle-timeout

Syntax **idle-timeout** *timeout*
no idle-timeout

Context config>subscr-mgmt>sla-prof>cat-map>category

Description This command defines the idle-timeout value.

The **no** form of this command reverts to the default.

Parameters *timeout* — Specifies the idle-timeout, in seconds

Values 60 to 15552000

idle-timeout-action

Syntax **idle-timeout-action** {**shcv-check** | **terminate**}
no idle-timeout-action

Context config>subscr-mgmt>sla-prof>cat-map>category

Description This command defines the action to be executed when the idle-timeout is reached. The action is performed for all hosts associated with the sla-profile instance.

The **no** form of this command reverts to the default.

Default idle-timeout-action terminate

Parameters **shcv-check** — Performs a subscriber host connectivity verification check (IPoE hosts only)

Note: Host connectivity verification must be enabled on the group-interface where the host is connected.

If the check is successful, the hosts are not disconnected and the idle-timeout timer is reset.

If the check fails, the hosts are deleted, similar as for **idle-timeout-action terminate**.

terminate — Deletes the subscriber host from the system: for PPP hosts, a terminate request is send; for IPoE hosts a DHCP release is send to the DHCP server

credit-type-override

Syntax **credit-type-override** {**volume** | **time**}
no credit-type-override

Context config>subscr-mgmt>cat-map>category

Description This command overrides the **credit-type** configured in the **config>subscr-mgmt>cat-map** context for the given category.

The **no** form of this command reverts to the default.

Parameters **volume** — If different than the value specified in the **credit-type** command, the value overrides the credit-type

time — If different than the value specified in the **credit-type** command, the value overrides the credit-type

default-credit

Syntax **default-credit volume** *credits* [**bytes** | **kilobytes** | **megabytes** | **gigabytes**]
default-credit time *seconds*
no default-credit

Context config>subscr-mgmt>cat-map>category

Description This command configures the default credit values for RADIUS credit control and Diameter Gy application credit control.

For RADIUS credit control, this command configures the default time or volume credit for this category. The default credit is used during initial setup when no quota is received from the RADIUS server.

For Diameter Gy credit control, this command specifies the interim credit assigned to this category (rating group) when Extended Failure Handling (EFH) is enabled. This command is ignored when EFH is disabled.

The **no** form of this command reverts to the default.

Parameters **volume** *credits* [**bytes** | **kilobytes** | **megabytes** | **gigabytes**] — Specifies the default value for the volume credit and the unit in which the default value is expressed

Values 1 to 4294967295 (minimum 1 byte)

time *seconds* — Specifies the default value for the time credit, in seconds

Values 1 to 4294967295 (minimum 1 second)

exhausted-credit-service-level

Syntax [**no**] **exhausted-credit-service-level**

Context config>subscr-mgmt>cat-map>category

Description This command enables the context to configure the exhausted credit service level.

The **no** form of this command reverts to the default.

Default exhausted-credit-service-level

egress-ip-filter-entries

Syntax [**no**] **egress-ip-filter-entries**

Context config>subscr-mgmt>cat-map>category>exh-lvl

Description This command enables the context to configure the egress IP filter parameters.

The **no** form of this command reverts to the default.

Default egress-ip-filter-entries

egress-ipv6-filter-entries

Syntax [**no**] **egress-ipv6-filter-entries**

Context config>subscr-mgmt>cat-map>category>exh-lvl

Description This command enables the context to configure the egress IPv6 filter parameters.

The **no** form of this command reverts to the default.

Default egress-ipv6-filter-entries

ingress-ip-filter-entries

Syntax [**no**] **ingress-ip-filter-entries**

Context config>subscr-mgmt>cat-map>category>exh-lvl

Description This command enables the context to configure the ingress IP filter parameters.

The **no** form of this command reverts to the default.

Default ingress-ip-filter-entries

ingress-ipv6-filter-entries

Syntax [**no**] **ingress-ipv6-filter-entries**

Context config>subscr-mgmt>cat-map>category>exh-lvl

Description This command enables the context to configure the ingress IPv6 filter parameters.

The **no** form of this command reverts to the default.

Default ingress-ipv6-filter-entries

pir

Syntax **pir** *pir-rate*
pir max
no pir

Context config>subscr-mgmt>cat-map>category>exh-lvl

Description This command configures the PIR which is enforced for all queues pertaining to this category.

The **no** form of this command reverts to the default.

Parameters *pir-rate* — Specifies the amount of bandwidth in kbytes/s (thousand bits per second)

Values 1 to 40000000

max — Specifies to use the maximum amount of bandwidth

entry

Syntax **entry** *entry-id* [**create**]
no entry *entry-id*

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip

config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6

Description This command configures the IP filter entry.

The **no** form of this command reverts to the default.

Parameters *entry-id* — Specifies the entry ID.

Values 1 to 65535

create — Keyword used to create an entry. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

action

Syntax **action drop**
action forward
action http-redirect *url* [**allow-override**]
no action

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry

Description This command configures the action for the filter entry.

The **no** form of this command reverts to the default.

Default action drop

Parameters **drop** — Specifies to drop the packets matching the IP filter entry.

forward — Specifies to forward the packets matching the IP filter entry.

http-redirect *url* [**allow-override**] — Specifies the HTTP web address, up to 255 characters, that is sent to the user's browser for redirection.

The specified URL can be overridden by a Diameter Credit Control Server when the following conditions are met:

- a Final-Unit-Indication AVP is present in the Multiple-Services-Credit-Control AVP of a CCA message
- the Final-Unit-Action AVP is set to REDIRECT (1)
- a Redirect-Server AVP is included with the following:
  - the Redirect-Address-Type AVP set to URL (2)
  - the Redirect-Server-Address AVP containing the URL to use for this rating group (**category-map**)

- the out of credit action for the corresponding rating group is set to **change-service-level** using one of the following commands:
  - **configure>subscriber-mgmt>credit-control-policy** *policy-name***>out-of-credit-action change-service-level**
  - **configure>subscriber-mgmt>category-map** *category-map-name* **category** *category-name* **>ut-of-credit-action-override change-service-level**
- an IPv4 HTTP redirect action with **allow-override** is specified in the exhausted credit service level context for the corresponding rating group using the command **configure>subscriber-mgmt>category-map** *category-map-name* **category** *category-name*>**exhausted-credit-service-level>ingress-ip-filter-entries**>**entry** *entry-id*>**action http-redirect** *url* **allow-override**

In all other cases, the URL specified in the Redirect-Server-Address AVP is ignored and the configured URL is used. The URL received from the Credit Control Server is included in the output of **show service active-subscribers credit-control**. The **allow-override** is ignored for RADIUS credit control.

The following variables can optionally be added in the configured URL (http-redirect url) and in the override URL from the Credit Control Server (Redirect-Server-Address AVP):

- $IP – Customer's IP address
- $MAC – Customer's MAC address
- $URL – Original requested URL
- $SAP – Customer's SAP
- $SUB – Customer's subscriber identification string
- $CID – string that represents the circuit-id or interface-id of the subscriber host (hexadecimal format)
- $RID – string that represents the remote-id of the subscriber host (hexadecimal format)

action

Syntax **action drop**
action forward
no action

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry

Description This command configures the action for the filter entry.

The **no** form of this command reverts to the default.

Default action drop

Parameters **drop** — Specifies to drop the packets matching the IP filter entry

forward — Specifies to forward the packets matching the IP filter entry

match

Syntax **match** [**protocol** *protocol-id*]
match [**next-header** *next-header*]
no match

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry

Description This command configures the match criteria for this IP filter entry.

The **no** form of this command reverts to the default.

Parameters *protocol-id* — Specifies the protocol ID or protocol name accepted in DHB

Values protocol-number — [0 to 255]D
[0x0 to 0xFF]H
[0b0 to 0b11111111]B

protocol-name — none, icmp, igmp, ip, tcp, egp, igp, udp, rdp, ipv6, ipv6-route, ipv6-frag, idrp, rsvp, gre, ipv6-icmp, ipv6-no-nxt, ipv6-opts, iso-ip, eigrp, ospf-igp, ether-ip, encap, pnni, pim, vrrp, l2tp, stp, ptp, isis, crtp, crudp, sctp

* - udp/tcp wildcard

next-header — Specifies the protocol numbers accepted in DHB for IPv6 filter entries

Values protocol-number — [1 to 42 | 45 to 49 | 52 to 59 | 61 to 255]D

[0x0 to 0x2A, 0x2D to 0x31, 0x34 to 0x3B, 0x3D to 0xFF]H [0b0 to 0b101010, 0b101101 to 0b110001, 0b110100 to 0b111011, 0b111101 to 0b11111111]B

protocol-name — none, icmp, igmp, ip, tcp, egp, igp, udp, rdp, ipv6, idrp, rsvp, gre, ipv6-icmp, ipv6-no-nxt, iso-ip, eigrp, ospf-igp, ether-ip, encap, pnni, pim, vrrp, l2tp, stp, ptp, isis, crtp, crudp, sctp

* - udp/tcp wildcard

dscp

Syntax **dscp** *dscp-name*
no dscp

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match

config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match

Description This command configures DSCP match conditions.

The **no** form of this command reverts to the default.

Parameters *dscp-name* — Specifies the DSCP name, up to 32 characters

dst-ip

Syntax **dst-ip** {*ipv6-address* | *prefix-length*}
dst-ip {*ip-address/mask* | *ip-address netmask*
no dst-ip

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match

Description This command configures the destination IP match condition.

The **no** form of this command reverts to the default.

Parameters *ipv6-address* — Specifies the IPv6 address (applies only to the 7750 SR)

Values

ipv6-address	x:x:x:x:x:x:x:x (where x is [0 to FFFFH])
	:x:x:x:x:x:d.d.d.d (where d is [0 to 255]D)

prefix-length — Specifies the prefix length (applies only to the 7750 SR)

Values 1 to 128

dst-port

Syntax **dst-port** {**lt** | **gt** | **eq**} *dst-port-number*
dst-port range *start end*
no dst-port

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match

Description This command configures the destination port match condition.

The **no** form of this command reverts to the default.

Parameters **lt** | **gt** | **eq** — Specifies the operator.

dst-port-number — Specifies the destination port number as a decimal hex or binary

Values 0 to 65535

fragment

Syntax **fragment** {**true** | **false**}

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match

Description This command configures the fragmentation match condition.

The **no** form of this command reverts to the default.

Parameters **true** — Enables fragmentation matching

false — Disables fragmentation matching

icmp-code

Syntax **icmp-code** *icmp-code*
no icmp-code

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match

Description This command configures the ICMP code match condition.

The **no** form of this command reverts to the default.

Parameters *icmp-code* — Specifies the ICMP code numbers accepted in DHB

Values 0 to 255

icmp-type

Syntax **icmp-type** *icmp-type*
no icmp-type

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match

Description This command configures the ICMP type match condition.

The **no** form of this command reverts to the default.

Parameters *icmp-type* — Specifies the ICMP type numbers accepted in DHB

Values 0 to 255

ip-option

Syntax **ip-option** *ip-option-value* [*ip-option-mask*]
no ip-option

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match

Description This command configures the IP option match condition.

The **no** form of this command reverts to the default.

Parameters *ip-option-value* — Specifies the IP option value as a decimal hex or binary

Values 0 to 255

ip-option-mask — Specifies the IP option mask as a decimal hex or binary

Values 0 to 255

multiple-option

Syntax **multiple-option** {**true** | **false**}

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match

Description This command configures the multiple-option match condition.

The **no** form of this command reverts to the default.

Parameters **true** — Enables checking the number of IP options in the IP header

false — Disables checking the number of IP options in the IP header

option-present

Syntax **option-present** {**true** | **false**}

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match

Description This command configures the option-present match condition.

The **no** form of this command reverts to the default.

Parameters **true** — Enables checking for the presence of IP options in the IP header

false — Disables checking for the presence of IP options in the IP header

src-ip

Syntax **src-ip** {*ip-address/mask* | *ip-address netmask*}
src-ip {*ipv6-address* | *prefix-length*}
no src-ip

Context config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match

Description This command configures the source IP match condition.

The **no** form of this command reverts to the default.

Parameters *ipv6-address* — Specifies the IPv6 address (applies only to the 7750 SR)

Values

ipv6-address	x:x:x:x:x:x:x:x (where x is [0 to FFFFH])
	x:x:x:x:x:x:d.d.d.d (where d is [0 to 255]D)

prefix-length — Specifies the prefix length (applies only to the 7750 SR)

Values 1 to 128

ip-address/mask — Specifies the IPv4 address and mask

Values ip-address a.b.c.d
mask 0 to 32

netmask — Specifies the mask, expressed as a dotted quad

Values a.b.c.d

src-port

Syntax **src-port** {**lt** | **gt** | **eq**} *src-port-number*
src-port range *start end*
no src-port

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match

Description This command configures the source port match condition.

The **no** form of this command reverts to the default.

Parameters **lt | gt | eq** — Specifies the operators.

src-port-number — Specifies the source port number as a decimal hex or binary

Values 0 to 65535

dst-port-number — Specifies the destination port number as a decimal hex or binary

Values 0 to 65535

tcp-ack

Syntax **tcp-ack** {**true** | **false**}
no tcp-ack

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match

Description This command configures the TCP ACK match condition.

The **no** form of this command reverts to the default.

Parameters **true** — Enables checking for the ACK bit

false — Disables checking for the ACK bit

tcp-syn

Syntax **tcp-syn** {**true** | **false**}
no tcp-syn

Context config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>ingr-ipv6>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ip>entry>match
config>subscr-mgmt>cat-map>category>exh-lvl>egr-ipv6>entry>match

Description This command configures the TCP SYN match condition.

The **no** form of this command reverts to the default.

Parameters **true** — Enables checking for the SYN bit

false — Disables checking for the SYN bit

pir

Syntax **pir** *pir-rate*
pir max
no pir

Context config>subscr-mgmt>cat-map>category>svc-lvl

Description This command configures the PIR which is enforced for all queues pertaining to this category.

Parameters *pir-rate* — Specifies the amount of bandwidth in kilobits per second (thousand bits per second)

Values 1 to 40000000

max — Specifies to use the maximum amount of bandwidth

out-of-credit-action-override

Syntax **out-of-credit-action-override** {**continue** | **block-category** | **change-service-level**}
no out-of-credit-action-override

Context config>subscr-mgmt>cat-map>category

Description This command specifies the action to be taken if the credit is exhausted.

Parameters **continue** — Specifies to continue when running out of credit

block-category — Specifies to block the category when running out of credit

change-service-level — Specifies to change the service level when running out of credit

policer

Syntax **policer** *policer-id* {**ingress-only** | **egress-only** | **ingress-egress**}
no policer *policer-id*

Context config>subscr-mgmt>cat-map>category

Description This command configures a policer in this category.

The **no** form of this command reverts to the default.

Parameters *policer-id* — Specifies an existing policer identifier

Values 1 to 63

ingress-only — Specifies that ingress policers are defined in this category

egress-only — Specifies that egress policers are defined in this category

ingress-egress — Specifies that ingress and egress policers are defined in this category

queue

Syntax **queue** *queue-id* {**ingress-only** | **egress-only** | **ingress-egress**}
no queue *queue-id*

Context config>subscr-mgmt>cat-map>category

Description This command configures a queue in this category.

The **no** form of this command reverts to the default.

Parameters *queue-id* — Specifies the queue ID for this instances. Each queue nominated in the category map is monitored for activity (over a period of approximately 60 seconds), should the activity fall below the threshold value then a time is started. Whenever this timer exceeds the configured timeout under the idle-timeout the action (currently disconnect) is executed for that subscriber and all hosts under that given SLA-profile-instance.

Values 1 to 32

ingress-only — Specifies that ingress queues are defined in this category

egress-only — Specifies that egress queues are defined in this category

ingress-egress — Specifies that ingress and egress queues are defined in this category

rating-group

Syntax **rating-group** *rating-group-id*
no rating-group

Context config>subscr-mgmt>cat-map>category

Description This command configures the rating group applicable for this category.

The **no** form of this command reverts to the default.

Parameters *rating-group-id* — Specifies the rating group applicable for this category

credit-exhaust-threshold

Syntax **credit-exhaust-threshold** *threshold-percentage*
no credit-exhaust-threshold

Context config>subscr-mgmt>cat-map

Description This command specifies the credit exhaust threshold taken into account to take action.

The **no** form of this command reverts the configured value to the default.

Default credit-exhaust-threshold 100

Parameters *threshold-percentage* — Specifies the percentage to use for the credit exhaust threshold

Values 50 to 100

credit-type

Syntax **credit-type** {**volume** | **time**}
no credit-type

Context config>subscr-mgmt>cat-map

Description This command specifies whether volume or time based accounting is performed.

The **no** form of this command reverts to the default.

Default credit-type volume

Parameters **volume** — Specifies volume-based accounting

time — Specifies time-based accounting

gx-session-level-usage

Syntax [**no**] **gx-session-level-usage**

Context config>subscr-mgmt>cat-map

Description This command controls the instantiation of an internal category required for Diameter Gx session level Usage Monitoring (per IP-CAN session).

When configured, Gx session level Usage Monitoring can be enabled for sessions associated with this category map.

The internal category for Gx session level Usage Monitoring is counted against the maximum of sixteen categories that can be configured.

When not configured (default), then no internal category is instantiated and Gx session level Usage Monitoring cannot be enabled for sessions associated with this category map.

9.25.2.1.7 Diameter Peer Policy Commands

diameter-peer-policy

Syntax **diameter-peer-policy** *peer-policy-name* [**role** {**client** | **proxy**}] [**create**]
no diameter-peer-policy

Context config>aaa

Description This command creates a base diameter policy with up to 5 peers. There is a (TCP) connection created to each peer while only two peers can be active (used by applications) simultaneously. Various diameter applications can reference this policy.

The **no** form of this command reverts to the default.

Parameters *peer-policy-name* — Specifies the name of the policy that is created

client — Specifies that diameter is configured as client. The client initiate peering connections towards the server or Diameter proxy. Various applications such as Gx, Gy or NASREQ are layered directly on top of the Diameter client.

proxy — Specifies that diameter is configured as proxy. Diameter proxy is used to provide multi-chassis redundancy and it can assumes active or standby state. The proxy relays messages between the Diameter client on one side and the server on the other side.

create — Specifies the keyword used to create the **diameter-peer-policy**. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

diameter-application-policy

Syntax **diameter-application-policy** *application-policy-name* [**create**]
no diameter-application-policy *application-policy-name*

Context config>subscr-mgmt
config>service>ies>sub-if>grp-if

Description This command enables the context do create and configure diameter application policy.

The **no** form of this command reverts to the default.

Parameters *application-policy-name* — Specifies the name of the diameter policy, up to 32 characters

create — Specifies the keyword used to create a diameter application policy. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

diameter-node

Syntax **diameter-node** *origin-host-string* **destination-realm** *destination-realm-string*
no diameter-node

Context config>subscr-mgmt>diam-app-plcy

Description This command configures the Diameter node for this Diameter application policy.

diameter-peer-policy

Syntax **diameter-peer-policy** *referenced-policy-name*
no diameter-peer-policy

Context config>subscr-mgmt>diam-app-plcy
config>service>ies>sub-if>grp-if

Description This command is used by an application (DCCA, Gx, policy-management application, and so on) to reference a base diameter peer policy that the application will use.

The **no** form of this command reverts to the default.

Parameters *referenced-policy-name* — Specifies the name of the referenced policy

applications

Syntax **applications** {[**gx**] [**gy**] [**nasreq**]}
no applications

Context config>aaa>diam-peer-plcy

Description This command specifies which applications are advertised in the Capability Exchange Request (CER) messages sent on the peers.

Applications that can be configured on a Diameter peer policy:

- client and proxy role:
  - gx
  - nasreq
  - gx nasreq
- client role only:
  - gy

Note: Gx and nasreq applications can be enabled simultaneously on a single diameter peer.

The **no** form of this command reverts to the default.

Parameters **gx** — Specifies that Gx application support is advertised in CER

gy — Specifies that Gy (DCCA) application support is advertised in CER

nasreq — Specifies that NASREQ application support is advertised in CER

application

Syntax **application** {**gx** | **gy** | **nasreq**}
no application

Context config>subscr-mgmt>diam-appl-plcy

Description This command specifies the Diameter application for which this policy contains the configuration details, such as AVPs to include and their format.

Applications are mutually exclusive.

The **no** form of this command reverts to the default.

Parameters **gx** — Specifies that Gx is the supported application of this DIAMETER policy

gy — Specifies that Gy is the supported application of this DIAMETER policy

nasreq — .Specifies that NASREQ is the supported application of this DIAMETER policy

ipv6-source-address

Syntax **ipv6-source-address** *ipv6-address*
no ipv6-source-address

Context config>aaa>diam-peer-plcy
config>aaa>diam>node

Description This command configures IPv6 source address that the SR OS node will use for its peering connection.

The **no** form of the command removes the IPv6 source address from the configuration.

Parameters *ipv6-address* — Specifies the source IPv6 address to use for outgoing diameter messages to an IPv6-reachable peer

Values

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x - [0 to FFFF]H
d - [0 to 255]D

origin-host

Syntax **origin-host** *origin-host-string*
no origin-host

Context config>aaa>diam-peer-plcy

Description This command configures the Origin-Host AVP that will be sent in all Diameter messages. Together with the Origin-Realm AVP, these two AVPs form a Diameter Identity.

This command is applicable only to legacy implementations of Diameter base in the SR OS.

The **no** form of this command removes the origin host string from the configuration.

Parameters *origin-host-string* — Specifies the Origin-Host AVP (Attribute Value Pair) used by this policy, up to 80 characters

origin-realm

Syntax **origin-realm** *origin-realm-string*
no origin-realm

Context config>aaa>diam-peer-plcy
config>aaa>diam-peer-plcy>peer

Description This command configures the Origin-Realm AVP that will be sent in all Diameter messages. Together with the Origin-Host AVP, these two AVPs form a Diameter Identity.

This command is applicable only to legacy implementations of Diameter base in the SR OS.

The **no** form of this command reverts to the default.

Parameters *origin-realm-string* — Specifies the origin-realm AVP (Attribute Value Pair) used by this policy, up to 80 characters

peer

Syntax **peer** *name* [**create**]
no peer name

Context config>aaa>diam-peer-plcy

Description This command creates context for diameter peer configuration. Up to five Diameter peers can be configured inside a **diameter-peer-policy**. Only the two peers with the highest preference in the peer table are considered for use (the primary and secondary). Other peers can be the Open state and merely run keepalives (watchdog-request or answer messages). After the primary peer fails, the secondary peer will be used as long as the last transaction on it has succeeded (stickiness). Another peer in the Open state will become secondary. Load balancing between peers is not supported.

This command is applicable only to legacy implementations of Diameter base in the SR OS.

The **no** form of this command reverts to the default.

Parameters *name* — Specifies the peer name, up to of 32 characters

create — Keyword used to create the peer. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

address

Syntax **address** *ip-address*
no address

Context config>aaa>diam-peer-plcy>peer

Description This command configures the IPv4 or IPv6 address of the diameter peer.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the IPv4 or IPv6 address of the diameter peer

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D

destination-host

Syntax **destination-host** *destination-host-string*
no destination-host

Context config>aaa>diam-peer-plcy>peer

Description This command configures the Destination-Host AVP that will be sent in CCR-I/U/T and RAA messages. If the destination host name is not explicitly set by configuration, it will be learned from CCA or RAR messages. In other words, the origin host name received in the CCA or RAR message will be used to populate or replace the destination host name for a session in SR OS node.

This command is applicable only to legacy implementations of Diameter base in the SR OS.

The **no** form of this command reverts to the default.

Parameters *destination-host-string* — Specifies the destination host name, up to 80 characters

preference

Syntax **preference** *preference*
no preference

Context config>aaa>diam-peer-plcy>peer

Description This command configures preference for a peer. Only the two peers with the highest preference will be used to carry traffic. Other peers can be in the Open state waiting to be utilized in case that one of the two active peers fail. Once the primary peer fails, the secondary peer will be used to carry traffic as long as the last transaction on it has succeeded (stickiness). Another peer in the Open state will become secondary. Load balancing between peers is not supported.

This command is applicable only to legacy implementations of Diameter base in the SR OS.

The **no** form of this command reverts to the default value.

Parameters *preference* — Specifies the preference of this diameter policy peer

Values 1 to 100

transaction-timer

Syntax **transaction-timer** *seconds*
no transaction-timer

Context config>aaa>diam-peer-plcy
config>aaa>diam-peer-plcy>peer

Description This command defines the time-out value for the Base Diameter messages (CER, DPR). When the transaction-timer expires, an appropriate action is taken for each message type.

This timer is used in the following cases:

- Opening the TCP connection (and completing the 3-way handshake) - if the TCP ACK is not received within the time specified by the transaction-timer, the TCP connection is closed and the connection-timer is started waiting for the new connection to be initiated
- Capability Exchange — If the response to the CER message (CEA) is not received within the time specified by the transaction-timer, the peer connection is closed and the connection-timer is started waiting for the new connection to be initiated.
- Peer disconnect Request — If the response to the DPR message is not received (DPA) within the time specified by the **transaction-timer**, the peer connection is closed.

This command is applicable only to legacy implementations of Diameter base in the SR OS.

The **no** form of this command reverts to the default.

Parameters *seconds* — Specifies the policy peer transaction timer, in seconds

Values 1 to 1000

transport

Syntax **transport tcp port** *port*
no transport

Context config>aaa>diam-peer-plcy>peer

Description This command defines the destination TCP port of the connection channel.

This command is applicable only to legacy implementations of Diameter base in the SR OS.

The **no** form of this command reverts to the default.

Default transport tcp port 3868

Parameters *port* — Specifies the transport protocol port number used toward this policy peer

Values 1 to 65535

destination-realm

Syntax **destination-realm** *destination-realm-string*
no destination-realm

Context config>aaa>diam-peer-plcy>peer

Description This command configures the destination-realm AVP that will be sent in CCR-I/U/T and RAA messages. The destination realm name cannot be learned dynamically from the CCA or RAR messages and therefore it should be explicitly configured in SR OS node. Once configured, it cannot be changed while peers are open.

This command is applicable only to legacy implementations of Diameter base in the SR OS.

The **no** form of this command reverts to the default.

Parameters *destination-realm-string* — Specifies the destination realm name, up to 80 characters

watchdog-timer

Syntax **watchdog-timer** *seconds*
no watchdog-timer

Context config>aaa>diam-peer-plcy
config>aaa>diam-peer-plcy>peer

Description This command configures the interval between consecutive watchdog messages. On the first timeout of the DWR, the SR OS node will resend the DWR message. The peer is still operational during this time. On the second timeout, the peer will transition into a suspended mode and the peer-failover procedure will be initiated (if the peer-failover is enabled by configuration). In this state the peer is not used for new transactions. At the same time, the cooldown procedure is started which means that it would take three successful DWR/DWA message exchanges to reinstate the peer in a fully operation state.

On the third timeout, the peer is removed and its connection is closed.

This command is applicable only to legacy implementations of Diameter base in the SR OS.

The **no** form of this command reverts to the default.

Default watchdog-timer 30

Parameters *seconds* — Specifies the device watchdog timer, in seconds, used by this policy peer

Values 1 to 1000

proxy

Syntax **proxy**

Context config>aaa>diam-peer-plcy

Description This command enables the context to configure the diameter peer policy proxy.

mcs-peer

Syntax **mcs-peer** *ip-address* **sync-tag** *name*
no mcs-peer

Context config>aaa>diam-peer-plcy>proxy

Description This command configures the IP address and the synchronization tag for the multi-chassis synchronization (MCS) peer.

This command is applicable only to legacy implementations of Diameter base in SR OS.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the IPv4 or IPv6 address of the MCS peer

name — Specifies the sync tag name, up to 32 characters. The name must match on both nodes for the pair of redundant diameter proxies.

router

Syntax **router** *router-instance* **address** *ip-address*
no router

Context config>aaa>diam-peer-plcy>proxy

Description This command configures the IPv4or IPv6 address and routing context on which the Diameter proxy is listening for the client connections and is applicable to Diameter proxy only.

This command is applicable only to legacy implementations of Diameter base in the SR OS.

The **no** form of this command reverts to the default.

Parameters *router-instance* — Specifies the routing context in which Diameter proxy is configured

ip-address — Specifies the IPv4 or IPv6 address to which the Diameter proxy is listening

Values

ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D

python-policy

Syntax **python-policy** [*python-policy-name*]
no python-policy

Context config>aaa>diam-peer-plcy
config>aaa>diam>node

Description This command references python-policy name that will be used for diameter processing.

The **no** form of this command reverts to the default.

Parameters *python-policy-name* — Specifies the name of the Python policy, up to 32 characters

router

Syntax **router** *router-instance*
router service *service-name*
no router

Context config>aaa>diam-peer-plcy

Description This command references the routing instance from which diameter peering is instantiated.

The **no** form of this command reverts to the default.

Parameters *router-instance* — Specify one of the following parameters for the router instance.

router-name — Specifies a router name, up to 32 characters, to be used in the match criteria.

Values Base, management

Default Base

service-id — Specifies an existing service ID to be used in the match criteria

Values 1 to 2147483647

service-name service-name — Specifies an existing service name, up to 64 characters

source-address

Syntax **source-address** *ip-address*
no source-address

Context config>aaa>diam-peer-plcy
config>aaa>diam>node

Description This command configures IPv4 source address that the SR OS node will use for its peering connection.

The **no** form of this command removes the IP address from the configuration.

Parameters *ip-address* — Specifies IP prefix for the IP match criterion in dotted decimal notation

Values 0.0.0.0 **to** 255.255.255.255

vendor-support

Syntax **vendor-support** [**three-gpp** | **vodafone**]
no vendor-support

Context config>subscr-mgmt>diam-appl-plcy>gy
config>aaa>diam-peer-plcy

Description In a diameter peer policy, this command specifies the vendor support announced in the capability exchange. In a Gy diameter application policy, this command specifies the vendor specific attributes for the user sessions.

The **no** form of this command reverts to the default value.

Default vendor-support three-gpp

Parameters **three-gpp** — Specifies the 3GPP diameter policy vendor type

vodafone — Specifies the vodafone diameter policy vendor type

extended-failure-handling

Syntax **extended-failure-handling**

Context config>subscr-mgmt>diam-appl-plcy>gy

Description This command enables the context to configure Extended Failure Handling (EFH), a mechanism to establish a new Diameter Gy session with the Online Charging Server (OCS) after Credit Control Failure Handling (CCFH) CONTINUE is triggered.

interim-credit

Syntax **interim-credit**

Context config>subscr-mgmt>diam-appl-plcy>gy>efh

Description This command enables the context to configure interim credit parameters for Extended Failure Handling (EFH).

max-attempts

Syntax **max-attempts** *count*
max-attempts infinite
no max-attempts

Context config>subscr-mgmt>diam-appl-plcy>gy>efh>interim-c

Description This command configures the maximum number of attempts made to establish a new Diameter Gy session with the Online Charging Server (OCS) when Extended Failure Handling (EFH) is active.

A new attempt is made when the volume or time interim credit of a rating group is consumed or when the validity time expires for a rating group.

When the maximum number of attempts is reached, the user session associated with the Diameter session is terminated (the corresponding subscriber hosts are deleted from the system).

The **no** form of this command resets the value to the default value.

Default max-attempts 10

Parameters *count* — Specifies the maximum number attempts that is made to establish a Diameter Gy session with the OCS when EFH is active

Values 1 to 4294967295

infinite — Specifies that an unlimited number of attempts is made to establish a Diameter Gy session with the OCS when EFH is active

reporting

Syntax [**no**] **reporting**

Context config>subscr-mgmt>diam-appl-plcy>gy>efh>interim-c

Description This command enables reporting of the used interim credit for each rating group when a new Diameter Gy session is successfully established with the Online Charging Server (OCS).

When enabled, the used interim credit report includes:

- the unreported used credit assigned via the initial Diameter Gy session when Extended Failure Handling (EFH) became active
- the used interim credits during EFH
- the used credits assigned via the new established Diameter Gy session

The **no** form of this command disables interim credit reporting. Only credit assigned via the new established Diameter Gy session is reported.

validity-time

Syntax **validity-time** *seconds*
no validity-time

Context config>subscr-mgmt>diam-appl-plcy>gy>efh>interim-c

Description This command configures the validity time for the interim credit allocated to rating groups of a Diameter Gy session when Extended Failure Handling (EFH) is active. When either the allocated interim credit is consumed or the validity time expires, a new attempt is made to establish a Diameter Gy session with the Online Charging Server (OCS).

The validity time applies to all interim credit allocated via the **config>subscr-mgmt>diam-appl-plcy** *application-policy-name*>**gy>extended-failure-handling>interim-credit>volume** and **config>subscr -mgmt>category-map** *category-map-name*>**category** *category-name*>**default-credit** CLI commands.

A validity time value of 0 (zero) disables the validity time for the assigned interim credit.

The **no** form of this command resets the value to the default value.

Default validity-time 1800

Parameters *seconds* — Specifies the validity time, in seconds, applicable to the interim credit

Values 0 to 4294967295

volume

Syntax **volume** *credits* {**bytes** | **kilobytes** | **megabytes** | **gigabytes**}
no volume

Context config>subscr-mgmt>diam-appl-plcy>gy>efh>interim-c

Description This command configures the default volume interim credit that is allocated to all rating groups of a Diameter Gy session when Extended Failure Handling (EFH) is active and for which no default credit is configured at the category map category level.

The **no** form of this command resets the value to the default value.

Default volume 500 megabytes

Parameters *credits* — Specifies the amount of volume credit that is allocated by default to all rating groups of a Diameter Gy session when EFH is active

Values 1 to 4294967295

bytes | **kilobytes** | **megabytes** | **gigabytes** — Specifies whether credits are in bytes, kilobytes, megabytes, or gigabytes

new-session-id

Syntax [**no**] **new-session-id**

Context config>subscr-mgmt>diam-appl-plcy>gy>efh

Description This command determines the Diameter session ID when Extended Failure Handling (EFH) is active and an attempt is made to establish a new Diameter Gy session with the Online Charging Server (OCS). An attempt to establish a new Diameter Gy session is made when the allocated interim credit is used or the validity time expires for a rating group of a Diameter Gy session. The first attempt always uses a new Diameter session ID. This command controls the behavior for each subsequent attempt. The behavior is as follows:

- no new-session-id (default) — The same Diameter session ID is used for each subsequent attempt.
- new-session-id — A new Diameter session ID is used for each attempt.

The **no** form of this command reverts to the default.

shutdown

Syntax [**no**] **shutdown**

Context config>subscr-mgmt>diam-appl-plcy>gy>efh

Description This command enables or disables Extended Failure Handling (EFH) for Diameter Gy sessions established using this Diameter application policy.

The **no** form of this command enables EFH for Diameter Gy sessions established using this Diameter application policy.

The **shutdown** command disables EFH for Diameter Gy sessions established using this Diameter application policy.

The **no** form of this command reverts to the default.

Default shutdown

include-avp

Syntax [**no**] **include-avp**

Context config>subscr-mgmt>diam-appl-plcy>gy
config>subscr-mgmt>diam-appl-plcy>gx
config>subscr-mgmt>diam-appl-plcy>nasreq

Description This command enables the context to configure AVPs and their format to be included in Diameter Gx, Gy, or NASREQ application messages. For full description each AVP, refer to the 7750 SR and VSR RADIUS Attributes Reference Guide.

AVP name:

- an-gw-address
- apn-ambr
- called-station-id
- calling-station-id
- charging-characteristics
- dynamic-address-flag
- ip-can-type

- logical-access-id
- nas-port
- nas-port-id
- nas-port-type
- pdn-connection-id
- physical-access-id
- rai
- rat-type
- selection-mode
- sgsn-mcc-mnc
- supported-features
- user-equipment-info
- user-location-info

an-gw-address

Syntax [**no**] **an-gw-address**

Context config>subscr-mgmt>diam-appl-plcy>gx>include-avp

Description This command configures the IPv4 address of the node.

The **no** form of this command reverts to the default.

apn-ambr

Syntax [**no**] **apn-ambr**

Context config>subscr-mgmt>diam-appl-plcy>gx>include-avp

Description This command enables the inclusion of the APN-Aggregate-Max-Bitrate-DL and APN-Aggregate-Max-Bitrate-UL AVPs inside the QoS-Information AVP, as signaled in the incoming GTP message.

The **no** form of this command disables the inclusion of the AVPs.

called-station-id

Syntax [**no**] **called-station-id**

Context config>subscr-mgmt>diam-appl-plcy>gx>include-avp
config>subscr-mgmt>diam-appl-plcy>nasreq>include-avp

Description This command configures the MAC address of AP in WiFi.

The **no** form of this command reverts to the default.

calling-station-id

Syntax **calling-station-id** [**type** {**llid** | **mac** | **remote-id** | **sap-id** | **sap-string**}]
no calling-station-id

Context config>subscr-mgmt>diam-appl-plcy>gx>include-avp
config>subscr-mgmt>diam-appl-plcy>nasreq>include-avp

Description This command includes the calling-station-id AVP in the specified format.

The **no** form of this command reverts to the default.

Parameters **type** — Specifies the format of the Calling-Station-ID AVP

Values **llid** — The logical link identifier (LLID) is the mapping from a physical to logical identification of a subscriber line and supplied by a RADIUS llid-serv

mac — Specifies that the MAC address is sent.

remote-id — Specifies that the remote ID is sent

sap-id — Specifies that the sap-id is sent

sap-string — Specifies that the value is the inserted value set at the SAP level. If no **calling-station-id** value is set at the SAP level, the **calling-station-id** attribute will not be sent.

circuit-id

Syntax [**no**] **circuit-id**

Context config>subscr-mgmt>diam-appl-plcy>nasreq>include-avp

Description This command includes the Agent-Circuit-Id AVP.

The **no** form of this command reverts to the default.

imei

Syntax [**no**] **imei**

Context config>subscr-mgmt>diam-appl-plcy>nasreq>include-avp

Description This command enables the inclusion of the IMEI AVP, as signaled in the incoming GTP setup message.

The **no** form of this command disables the inclusion of the AVP.

ip-can-type

Syntax [**no**] **ip-can-type**

Context config>subscr-mgmt>diam-appl-plcy>gx>include-avp

Description This command includes the ip-can-type.

The **no** form of this command reverts to the default.

logical-access-id

Syntax [**no**] **logical-access-id**

Context config>subscr-mgmt>diam-appl-plcy>gx>include-avp

Description This command includes the logical-access-id.

The **no** form of this command reverts to the default.

nas-port

Syntax **nas-port** *binary-spec*
no nas-port

Context config>subscr-mgmt>diam-appl-plcy>gx>include-avp
config>subscr-mgmt>diam-appl-plcy>nasreq>include-avp

Description This command specifies the format of the 32 bit string used as value for the Nas-Port AVP.

Parameters *binary-spec* — Specifies the NAS-Port AVP format

Values

binary-spec	<bit-specification> <binary-spec>
bit-specification	0 | 1 | <bit-origin>
bit-origin	*<number-of-bits><origin>
number-of-bits	1 to 32
origin	s | m | p | o | i | v | c
	s - slot number
	m - MDA number
	p - port number or lag-id
	o - outer VLAN ID
	i - inner VLAN ID
	v - ATM VPI
	c - ATM VCI

nas-port-id

Syntax **nas-port-id** [**prefix-type** {**none** | **user-string**}] [**prefix-string** *prefix-string*] [**suffix-type** {**circuit-id** | **none** | **remote-id** | **user-string**}] [**suffix-string** *suffix-string*]
no nas-port-id

Context config>subscr-mgmt>diam-appl-plcy>gx>include-avp
config>subscr-mgmt>diam-appl-plcy>nasreq>include-avp

Description This command includes the Nas-Port-Id AVP.

Parameters **prefix-type** — Specifies what type of prefix is added to the NAS-Port-Id attribute if included in Nas-Port-Id AVP messages

Values **none** — No prefix is added
user-string — Specifies the user configurable string to be added as prefix to the NAS-Port-Id attribute if included in DIAMETER Gx messages

prefix-string — Specifies the user configurable string up to 8 characters, to be added as a prefix

suffix-type} — Specifies the suffix to be added to the NAS-Port attribute NAS-Port AVP

Values **none** — No suffix is added
circuit-id — Specifies the circuit-id is added as suffix-string
remote-id — Specifies the remote-id is added as suffix-string
user-string — Specifies a user configurable suffix-string is added

suffix-string — Specifies the string, up to 64 characters, to be added as suffix

nas-port-type

Syntax **nas-port-type**
nas-port-type [*type*]
no nas-port-type

Context config>subscr-mgmt>diam-appl-plcy>gx>include-avp
config>subscr-mgmt>diam-appl-plcy>nasreq>include-avp

Description This command includes the Nas-Port-Type AVP.

Parameters **none** — Specifies values as defined in RFC 2865, Remote Authentication Dial In User Service (RADIUS), and RFC 4603, Additional Values for the NAS-Port-Type Attribute.

type — Specifies the integer value for the Nas-Port-Type AVP.

Values 0 to 255

pdn-connection-id

Syntax [**no**] **pdn-connection-id**

Context config>subscr-mgmt>diam-appl-plcy>gx>include-avp

Description This command enables the inclusion of the PDN-Connection-Id AVP, which contains the APN as signaled in the incoming GTP setup message.

The **no** form of this command disables the inclusion of the AVP.

rai

Syntax [**no**] **rai**

Context config>subscr-mgmt>diam-appl-plcy>gx>include-avp

Description This command enables the inclusion of the RAI AVP as signaled in the incoming GTP setup message.

The **no** form of this command disables the inclusion of the AVP.

remote-id

Syntax [**no**] **remote-id**

Context config>subscr-mgmt>diam-appl-plcy>gx>include-avp
config>subscr-mgmt>diam-appl-plcy>nasreq>include-avp

Description This command enables the generation of the agent-remote-id for RADIUS.

The **no** form of this command reverts to the default.

physical-access-id

Syntax [**no**] **physical-access-id**

Context config>subscr-mgmt>diam-appl-plcy>gx>include-avp

Description This command includes the physical access ID.

The **no** form of this command reverts to the default.

rat-type

Syntax [**no**] **rat-type**

Context config>subscr-mgmt>diam-appl-plcy>gx>include-avp
config>subscr-mgmt>diam-appl-plcy>nasreq>include-avp

Description This command includes the RAT type.

The **no** form of this command reverts to the default.

sgsn-mcc-mnc

Syntax [**no**] **sgsn-mcc-mnc**

Context config>subscr-mgmt>diam-appl-plcy>gx>include-avp

Description This command enables the inclusion of the 3GPP-SGSN-MCC-MNC AVP, which contains the MCC and MNC as configured under **configure subscriber-mgmt gtp serving-network**.

The **no** form of this command disables the inclusion of the AVP

supported-features

Syntax [**no**] **supported-features**

Context config>subscr-mgmt>diam-appl-plcy>gx>include-avp

Description This command includes the **supported-features** in CCR messages.

The **no** form of this command resets the command to the default setting.

user-equipment-info

Syntax **user-equipment-info** [**type** *ue-info-type*]
no user-equipment-info

Context config>subscr-mgmt>diam-appl-plcy>gx>include-avp
config>subscr-mgmt>diam-appl-plcy>gy>include-avp

Description This command includes the **user-equipment-info** in CCR messages.

The **no** form of this command resets the command to the default setting.

Default user-equipment-info type mac

Parameters *ue-into-type* — Specifies what is included in the Diameter User-Equipment-Info attribute if included in Diameter Gx messages.

Values eui64, imeisv, mac, modified-eui64

user-location-info

Syntax [**no**] **user-location-info**

Context config>subscr-mgmt>diam-appl-plcy>gx>include-avp
config>subscr-mgmt>diam-appl-plcy>nasreq>include-avp

Description This command enables the inclusion of the 3GPP-User-Location-Information AVP as signaled in the incoming GTP setup message.

The **no** form of this command disables the inclusion of the AVP.

mac-format

Syntax **mac-format** *mac-format*
no mac-format

Context config>subscr-mgmt>diam-appl-plcy>gx
config>subscr-mgmt>diam-appl-plcy>nasreq
config>subscr-mgmt>diam-appl-plcy>gy

Description This command configures the format of the MAC address when reported in Gx, Gy, or NASREQ application message AVPs such as Calling-Station-Id or User-Name.

The **no** form of this command resets the command to the default setting.

Default mac-format ab

Parameters *mac-format* — Specifies the MAC address format.

Values
like ab: for 00:0c:f1:99:85:b8
or XY- for 00-0C-F1-99-85-B8
or mmmm. for 0002.03aa.abff
or xx for 000cf19985b8

report-ip-address-event

Syntax [**no**] **report-ip-address-event**

Context config>subscr-mgmt>diam-appl-plcy>gx

Description This command enables triggered CCR-u messages based on IP address allocation or de-allocation for the subscriber host.

If the requests for both IP address families (IPv4 and IPv6) arrive at approximately the same time, a single CCR-i is sent containing the IP addresses from both address families, IPv4 and IPv6 (NA, PD, or SLAAC). When the requests for IP addresses are not nearly simultaneous, the CCR-i contains only the IP address that was allocated first (the one that triggered the session creation). The request for the second IP address family, depending on configuration, triggers an additional CCR-u that will carry the IP address allocation update to the PCRF along with the UE_IP_ADDRESS_ALLOCATE (18) event. The CCR-u content mirrors the content of the CCR-i with the exception of already allocated IP address(es).

When this command is disabled, IP address-triggered CCR-u messages are not sent.

The **no** form of this command disables the command.

Default report-ip-address-event

3gpp-ggsn-address

Syntax [**no**] **3gpp-ggsn-address**

Context config>subscr-mgmt>diam-appl-plcy>gy>include-avp

Description This command includes the 3GPP GGSN Address AVP in the Diameter Gy messages. The value is set to the source IPv4 address that is used for outgoing diameter messages.

The **no** form of this command disables the command.

Default 3gpp-ggsn-address

3gpp-ggsn-ipv6-address

Syntax **3gpp-ggsn-ipv6-address**
no 3gpp-ggsn-ipv6-address

Context config>subscr-mgmt>diam-appl-plcy>gy>include-avp

Description This command includes the 3GPP GGSN IPv6 Address AVP in the Diameter Gy messages. It's value is set to the source IPv6 address that is used for outgoing diameter messages.

The **no** form of this command reverts to the default.

3gpp-charging-characteristics

Syntax [**no**] **3gpp-charging-characteristics**

Context config>subscr-mgmt>diam-appl-plcy>gy>avp

Description This command includes the 3GPP-Charging-Characteristics AVP in all Diameter DCCA CCR messages.

The **no** form of this command removes the 3GPP-Charging-Characteristics AVP from the Diameter DCCA CCR messages.

Default 3gpp-charging-characteristics

3gpp-charging-id

Syntax **3gpp-charging-id** {**auto** | **esm-info** | **id**}
no 3gpp-charging-id

Context config>subscr-mgmt>diam-appl-plcy>gy>avp

Description This command includes the 3GPP-Charging-Id AVP in all Diameter DCCA CCR messages. The value can be configured to contain a unique 32-bit identifier or ESM information.

The **no** form of this command removes the 3GPP-Charging-Id AVP from the Diameter DCCA CCR messages.

Default 3gpp-charging-id auto

Parameters **auto** — When the vendor support is 3GPP, this parameter specifies the AVP value for the *subscriber-id*, *sap-id*, and *sla-profile*. The default parameter value selected is **auto**.

esm-info — Specifies the AVP value to *subscriber-id*, *sap-id*, and *sla-profile*.

id — Specifies the AVP value asa unique 32-bit identifier per Diameter Gy session.

3gpp-gprs-negotiated-qos-profile

Syntax [**no**] **3gpp-gprs-negotiated-qos-profile**

Context config>subscr-mgmt>diam-appl-plcy>gy>avp

Description This command includes the 3GPP-GPRS-QoS-Negotiated-Profile AVP in all Diameter DCCA CCR messages. The value is the active SLA profile.

The **no** form of this command removes the 3GPP-GPRS-QoS-Negotiated-Profile AVP from the Diameter DCCA CCR messages.

Default 3gpp-gprs-negotiated-qos-profile

3gpp-imsi

Syntax **3gpp-imsi** {**circuit-id** | **imsi** | **subscriber-id**}

no 3gpp-imsi

Context config>subscr-mgmt>diam-appl-plcy>gy>avp

Description This command includes the 3GPP-IMSI AVP in all Diameter DCCA CCR messages.

The **no** form of this command removes the 3GPP-IMSI AVP from the Diameter DCCA CCR messages.

Parameters **circuit-id** — Specifies the circuit-id as DCCA IMSI AVP value.

subscriber-id — Specifies the subscriber-id as DCCA IMSI AVP value.

imsi — Specifies the IMSI as DCCA IMSI AVP value.

3gpp-nsapi

Syntax [**no**] **3gpp-nsapi**

Context config>subscr-mgmt>diam-appl-plcy>gy>avp

Description This command includes the 3GPP-NSAPI AVP in all Diameter DCCA CCR messages. The **3gpp-nsapi** avp value is created from a semi-colon separated concatenation of the system name, service ID, and SAP ID of the subscriber host or session.

The **no** form of this command removes the 3GPP-NSAPI AVP from the Diameter DCCA CCR messages.

3gpp-rat-type

Syntax **3gpp-rat-type** *value*
no 3gpp-rat-type

Context config>subscr-mgmt>diam-appl-plcy>gy>avp

Description This command includes the 3GPP-RAT-Type AVP with the specified value in all Diameter DCCA CCR messages.

The **no** form of this command removes the 3GPP-RAT-Type AVP from the Diameter DCCA CCR messages.

Parameters *value* — Specifies the 3GPP-RAT-Type AVP value to be included in Diameter DCCA CCR messages.

Values 1 to 255

3gpp-selection-mode

Syntax [**no**] **3gpp-selection-mode**

Context config>subscr-mgmt>diam-appl-plcy>gy>avp

Description This command includes the 3GPP-Selection-Mode AVP in all Diameter DCCA CCR messages.

The **no** form of this command removes the 3GPP-Selection-Mode AVP from the Diameter DCCA CCR messages.

3gpp-session-stop-indicator

Syntax [**no**] **3gpp-session-stop-indicator**

Context config>subscr-mgmt>diam-appl-plcy>gy>avp

Description This command includes the 3GPP-Session-Stop-Indicator AVP in all Diameter DCCA CCR-Terminate messages.

The **no** form of this command removes the 3GPP-Session-Stop-Indicator AVP from the Diameter DCCA CCR-Terminate messages.

3gpp-user-location-info

Syntax [**no**] **3gpp-user-location-info**

Context config>subscr-mgmt>diam-appl-plcy>gy>avp

Description This command, in case of ESM over GTP access, includes the 3GPP-User-Location-Info attribute. If **ps-information** is also enabled this is grouped within the PS information AVP.

The **no** version of this command disables inclusion of the ULI AVP.

Default 3gpp-user-location-info

address-avp

Syntax [**no**] **address-avp**

Context config>subscr-mgmt>diam-appl-plcy>gy>avp

Description This command includes the following subscriber host/session address/prefix AVPs in all Diameter DCCA CCR messages:

- [8] Framed-IP-Address

- [97] Framed-IPv6-Prefix
- [123] Delegated-IPv6-Prefix
- [6527-99] Alc-IPv6-Address

Note: Only the address/prefix of the subscriber host that triggered the creation of the Diameter Gy session is included.

The **no** form of this command removes the address AVPs from the Diameter DCCA CCR messages.

charging-rule-base-name

Syntax **charging-rule-base-name category-map-name**
charging-rule-base-name *string*
no charging-rule-base-name

Context config>subscr-mgmt>diam-appl-plcy>gy>avp

Description This command includes the Charging-Rule-Base-Name AVP with the specified value in all Diameter DCCA CCR messages.

The **no** form of this command removes the Charging-Rule-Base-Name AVP from the Diameter DCCA CCR messages.

Parameters **category-map-name** — This keyword specifies the name of the category-map in use.

string — Specifies a string of up to 64 characters.

ggsn-address

Syntax **ggsn-address** {**ipv4** | **ipv6**}
no ggsn-address

Context config>subscr-mgmt>diam-appl-plcy>gy>avp

Description The command includes the GGSN-Address AVP value in all Diameter DCCA CCR messages. The value is either the local IPv4 address or local IPv6 address used to set up the diameter peer.

The **no** form of this command removes the GGSN-Address AVP from the Diameter DCCA CCR messages.

Parameters **ipv4** | **ipv6** — Specifies to include either the IPv4 or IPv6 address.

pdp-context-type

Syntax [**no**] **pdp-context-type**

Context config>subscr-mgmt>diam-appl-plcy>gy>avp

Description This command includes the [3GPP-1247] PDP-Context-Type AVP in Diameter DCCA CCR-Initial messages.

The **no** form of this command removes the PDP-Context-Type AVP from the Diameter DCCA CCR-Initial messages.

ps-information

Syntax [**no**] **ps-information**

Context config>subscr-mgmt>diam-appl-plcy>gy>avp

Description This command includes the following AVPs in Diameter DCCA CCR messages encapsulated in [3GPP-873] Service-Information or [3GPP-874] PS-Information grouped AVPs:

- [8] Framed-IP-Address
- [30] Called-Station-Id
- [97] Framed-IPv6-Prefix
- [123] Delegated-IPv6-Prefix
- [6527-99] Alc-IPv6-Address
- [10415-1] 3GPP-IMSI
- [10415-2] 3GPP-Charging-Id
- [10415-5] 3GPP-GPRS-Negotiated-QoS-Profile
- [10415-7] 3GPP-GGSN-Address
- [10415-10] GGSN-NSAPI
- [10415-11] 3GPP-Session-Stop-Indicator
- [10415-12] 3GPP-Selection-Mode
- [10415-13] 3GPP-Charging-Characteristics
- [10415-16] 3GPP-GGSN-IPv6-Address
- [10415-21] 3GPP-RAT-Type
- [10415-847] GGSN-Address
- [10415-1004] Charging-Rule-Base-Name
- [10415-1247] PDP-Context-Type

The AVPs are included when configured in the **include-avp** context.

By default, these AVPs are included at the command level.

The **no** form of this command resets to the default setting.

called-station-id

Syntax **called-station-id** [*called-station-id*]
no called-station-id

Context config>subscr-mgmt>diam-appl-plcy>gy>include-avp

Description This command configures the value of the called station ID AVP.

The **no** form of this command returns the command to the default setting.

Parameters *called-station-id* — Specifies the called station ID up to 64 characters.

radius-user-name

Syntax [**no**] **radius-user-name**

Context config>subscr-mgmt>diam-appl-plcy>gy>include-avp

Description This command includes the RADIUS user name AVP in the Diameter Gy messages.

The **no** form of this command returns the command to the default setting.

service-context-id

Syntax **service-context-id** *name*
no service-context-id

Context config>subscr-mgmt>diam-appl-plcy>gy>include-avp

Description This command configures the value of the service context ID AVP.

The **no** form of this command returns the command to the default setting.

Parameters *name* — Specifies the service context ID AVP value up to 32 characters.

transaction-timer

Syntax [**no**] **transaction-timer** *transaction-timer*

Context config>aaa>diam-peer-plcy
config>aaa>diam-peer-plcy>peer

Description This command defines the timeout value for the Base Diameter messages (DWR, CER, DPR). Once the **transaction-timer** expires, an appropriate action is taken for each message type.

This timer is used in the following cases.

- Opening the TCP connection (and completing the three-way handshake). If the TCP ACK is not received within the time specified by the **transaction-timer**, the TCP connection is closed and the connection timer is started waiting for the new connection to be initiated.
- Capability Exchange — If the response to the CER message (CEA) is not received within the time specified by the **transaction-timer**, the peer connection is closed and the connection timer is started waiting for the new connection to be initiated.
- Peer disconnect Request — If the response to the DPR message is not received (DPA) within the time specified by the transaction-timer, the peer connection is closed.
- DWR Timeout — If the response to the DWR message is not received (DWA) within the time specified by the **transaction-timer**, the peer connection is not closed. Instead the peer will transition into a peer suspended mode and at the same time the watchdog timer is restarted.

The **no** form of this command returns the command to the default setting.

Default transaction-timer 30 seconds at diameter-base level

The default value at peer is taken from diameter-base.

Parameters *transaction-timer* — Specifies the Diameter peer policy transaction timer, in seconds.

Values 1 to 1000

Default 30 seconds at the diameter-base level of the CLI.
The default value in the peer context is taken from diameter-base.

gx

Syntax **gx**

Context config>subscr-mgmt>diam-appl-plcy

Description This command enables the context to configure Gx parameters.

gy

Syntax **gy**

Context config>subscr-mgmt>diam-appl-plcy

Description This command enables the context to configure Diameter Credit Control Application or Gy-specific options.

nasreq

Syntax **nasreq**

Context config>subscr-mgmt>diam-appl-plcy

Description This command enables the context to configure NASREQ application-specific attributes.

3gpp-qos-mapping

Syntax **3gpp-qos-mapping**

Context config>subscr-mgmt>diam-appl-plcy>gx

Description This command enables the configuration of mapping certain 3GPP-specific QoS attributes to SR OS objects.

apn-ambr-dl

Syntax **apn-ambr-dl agg-rate**
apn-ambr-dl arbiter *arbiter-name*
apn-ambr-dl hs-sla-agg-rate
apn-ambr-dl policer *policer-id*
apn-ambr-dl queue *queue-id*
apn-ambr-dl scheduler *scheduler-name*
no apn-ambr-dl

Context config>subscr-mgmt>diam-appl-plcy>gx>qos-map

Description This command configures the APN-Aggregate-Max-Bitrate-DL AVP. When enabled, the AVP is interpreted as a rate override for the specified egress QoS object. For queues and policers, the PIR is overridden.

This override uses the same QoS override mechanism as the native Gx and RADIUS-based QoS overrides. Therefore, a subsequent Gx/RADIUS-based override removes this override and an APN-AMBR based override removes any preceding Gx/RADIUS-based override.

The **no** form of this command disables the override mechanism based on APN-AMBR.

Parameters **agg-rate** — Specifies to map to an aggregate rate

arbiter-name — Specifies the name of the arbiter to be overridden

hs-sla-agg-rate — Specifies to map to an HS SLA aggregate rate

policer-id — Specifies the ID of the policer to be overridden

queue-id — Specifies the ID of the queue to be overridden

scheduler-name — Specifies the name of the scheduler to be overridden

apn-ambr-ul

Syntax **apn-ambr-ul arbiter** *arbiter-name*
apn-ambr-ul policer *policer-id*
apn-ambr-ul queue *queue-id*
apn-ambr-ul scheduler *scheduler-name*
no apn-ambr-ul

Context config>subscr-mgmt>diam-appl-plcy>gx>qos-map

Description This command configures the APN-Aggregate-Max-Bitrate-UL AVP. When enabled, the AVP is interpreted as a rate override for the specified egress QoS object. For queues and policers, the PIR is overridden.

This override uses the same QoS override mechanism as the native Gx and RADIUS-based QoS overrides. Therefore, a subsequent Gx/RADIUS-based override removes this override and an APN-AMBR based override removes any preceding Gx/RADIUS-based override.

The **no** form of this command disables the override mechanism based on APN-AMBR.

Parameters *arbiter-name* — Specifies the name of the arbiter to be overridden.

policer-id — Specifies the ID of the policer to be overridden.

queue-id — Specifies the ID of the queue to be overridden.

scheduler-name — Specifies the name of the scheduler to be overridden.

avp-subscription-id

Syntax **avp-subscription-id** *origin* [**type** *type*]
no avp-subscription-id

Context config>subscr-mgmt>diam-appl-plcy>gx
config>subscr-mgmt>diam-appl-plcy>gy

Description This command is used to provide identification information to the PCRF for the end user. Subscription-id is a grouped AVP. In case that parameter designated to be the subscription-id is not available, the subscription-avp will not be sent.

The **no** form of this command returns the command to the default setting.

Parameters *origin* — Specifies the origin of the information to send in the Subscription-Id data AVP.

Values **circuit-id** — Specifies the circuit ID. For Gy, the circuit id is an ASCII- converted value.

dual-stack-remote-id — Specifies the remote ID for IPv4 and IPv6. The enterprise-id field is stripped off from IPv6 remote ID. For Gy, the remote id is an ascii-converted value.

imei — Specifies the physical ID of the end device.

imsi — Specifies the SIM ID.

mac — Specifies the MAC address of the end device.

msisdn — Specifies the phone number of the end device.

nas-port-id — Specifies the NAS port ID which can be a prefix or suffix with a custom string to make it unique network wide.

subscriber-id — Specifies the subscriber ID.

username — Specifies that the username identifier can be of type **private** or **nai** which specifies the username is a PPP username (PAP/CHAP). When the PPP username is not available, the string in the Username attribute returned by RADIUS or NASREQ is used.

type — Specifies the type of the identifier stored in the Subscription-Id data AVP.

Values **e164** — Specifies the identifier is in international E.164 format (such as MSISDN).

imsi — Specifies the identifier is in international IMSI format according to the ITU-T E.212 numbering plan.

nai — Specifies the identifier is in the form of a Network Access Identifier as defined in RFC 2486.

private — Specifies the identifier is a private type identifier.

ccrt-replay

Syntax **ccrt-replay**

Context config>subscr-mgmt>diam-appl-plcy>gx
config>subscr-mgmt>diam-appl-plcy>gy

Description This command enables the context to configure CCR-terminate replay. CCR-t replay is enabled with a **no shutdown** of this context.

If a communication failure between client and server occurs, CCR-t replay enables the retransmission of CCR-t messages for a Gx or Gy session at a configured intervals until a valid response (CCA-t) is received or until the configured **max-lifetime** period expires, whichever comes first.

In Gx, replaying CCR-t messages ensures that the Gx session is cleared on the PCRF side in cases where the peering session to the PCRF was not available at the time that the initial and the first retransmitted CCR-t was sent.

In Gy, replaying CCR-t messages ensures that the final credit control usage reporting is not lost for billing by the OCS.

The subscriber host or session that triggered the Gx or Gy session that is in CCR-t replay mode is deleted from the system at the time that the initial CCR-t is sent. All resources associated with the subscriber host or session, such as queues, DHCP lease states, and PPPoE session states are released. The orphaned Gx and Gy sessions in replay mode are left in the system.

interval

Syntax **interval** *seconds*
no interval

Context config>subscr-mgmt>diam-appl-plcy>gx>ccrt-replay
config>subscr-mgmt>diam-appl-plcy>gy>ccrt-replay

Description This command specifies the interval at which CCR-t messages for Diameter Gx or Gy sessions that belong to the Diameter application policy are replayed, until a valid CCA-t response is received or until the configured **max-lifetime** period expires.

The **no** form of this command resets the interval to the default value of 3600 seconds.

Default no interval

Parameters *seconds* — Specifies the interval at which the CCR-t messages are replayed for a gx session. The messages are replayed until a valid CCA-t response is received or until a 24 hour period expires, whichever comes first.

Values 60 to 86400

Default 3600

max-lifetime

Syntax **max-lifetime** *hours*
no max-lifetime

Context config>subscr-mgmt>diam-appl-plcy>gx>ccrt-replay
config>subscr-mgmt>diam-appl-plcy>gy>ccrt-replay

Description This command specifies the maximum period of time that CCR-t messages for Diameter Gx or Gy sessions that belong to the Diameter application policy are replayed.

The **no** form of this command resets the maximum lifetime to the default value of 24 hours.

Default no max-lifetime

Parameters *hours* — Specifies the maximum lifetime after which the CCR-t messages are no longer replayed.

Values 1 to 24

Default 24

shutdown

Syntax [**no**] **shutdown**

Context config>subscr-mgmt>diam-appl-plcy>gx>ccrt-replay
config>subscr-mgmt>diam-appl-plcy>gy>ccrt-replay

Description This command, enables or disables the CCR-t replay function for all Gx or Gy sessions that belong to the diameter application policy. Sessions in ccrt replay are dropped when ccrt-replay is shutdown.

The **no** form of this command enables the CCR-t replay function.

Default shutdown

out-of-credit-reporting

Syntax **out-of-credit-reporting** {**final** | **quota-exhausted**}
no out-of-credit-reporting

Context config>subscr-mgmt>diam-appl-plcy>gy

Description This command changes the reporting reason in an intermediate interrogation when the final granted units have been consumed and a corresponding **out-of-credit-action** different from **disconnect-host** is started.

The **no** form of this command reverts to the default value.

Default out-of-credit-reporting final

Parameters **final** — Specifies the reporting reason in an intermediate interrogation when the final granted units have been consumed and a corresponding **out-of-credit-action** different from **disconnect-host** is started.

quota-exhausted — Specifies the reporting reason in an intermediate interrogation when the final granted units have been consumed and a corresponding **out-of-credit-action** different from **disconnect-host** is started.

on-failure

Syntax **on-failure** [**failover** {**enabled** | **disabled**}] [**handling** {**continue** | **retry-and-terminate** | **terminate**}]
no on-failure

Context config>subscr-mgmt>diam-appl-plcy

Description This command is applicable only to the R16.0R4 or later implementation of Diameter base in the SR OS.

This command configures session failure handling. The behavior of a Diameter application (Gx, Gy, or NASREQ) that fails to receive a response (an answer message CCA) to a transmitted request message (CCR) for a session, can be controlled by the Diameter server through two AVPs carried in CCA messages that are defined in RFC 4006, *Diameter Credit-Control Application*:

- CC-Session-Failover AVP
  - FAILOVER_NOT_SUPPORTED
  - FAILOVER_SUPPORTED
- Credit-Control-Failure-Handling AVP
  - TERMINATE
  - CONTINUE
  - RETRY_AND_TERMINATE

If those AVPs are not provided by the Diameter server, the local configuration provided by this command will take effect. This command defines the following:

- Retransmission behavior of the application in case that the response to a transmitted request message is not received. The Diameter base will notify the application after the lifetime of the transmitted request message expires (tx timeout) or the route to destination-realm is invalid. Depending on the configuration option, the application can then retransmit the message, or suppress retransmission.
- Handling behavior defines whether the session will be terminated or continue to exist in absence of the response from the Diameter server.

Default on-failure failover enabled handling terminate

Parameters **failover enabled** — Specifies that application session is allowed to retransmit the request message (CCR).

failover disabled — Specifies that the The application session is not allowed to retransmit the request message (CCR).

handling continue — Specifies that the sessions will continue to exist if the response to a transmitted CCR message is not received. Whether the transmitted message is re-transmitted depends on the failover configuration. In case of session initiation procedure in the Gx case (CCR-I timeout), the subscriber host is nonetheless instantiated with the default parameters, assuming that they are provided. If the default parameter are not provided, the subscriber host initiation fails.

handling retry-and-terminate — Specifies that the message is re-transmitted in case that the failover is enabled. If the response to the retransmitted request message is not received (re-transmission times out), the application session is terminated.

handling terminate — Specifies that the session is terminated if the response to the original request message is not received. No re-transmissions will be attempted, regardless of whether or not the failover is enabled.

on-failure (legacy)

Syntax **on-failure** [**failover** {**enabled** | **disabled**}] [**handling** {**continue** | **retry-and-terminate** | **terminate**}]
no on-failure

Context config>subscr-mgmt>diam-appl-plcy

Description This command is applicable only to legacy implementation of Diameter base in the SR OS.

This command configures session failure handling. The behavior of the application's session in case of a peer failure can be controlled by the Diameter server through two AVPs carried in CCA messages that are defined in RFC 4006, *Diameter Credit-Control Application*:

- CC-Session-Failover AVP
  - FAILOVER_NOT_SUPPORTED
  - FAILOVER_SUPPORTED
- Credit-Control-Failure-Handling AVP
  - TERMINATE
  - CONTINUE
  - RETRY_AND_TERMINATE

If those AVPs are not provided by the Diameter server, the local configuration provided by this command will take effect. This command defines the following:

- The peer-failover behavior to a secondary peer if the primary peer is unresponsive. From the application point of view, the primary peer is considered unresponsive if the request message that was sent to it, times out. The time out of the message is determined by the tx-timer command.

  The peer-failover action based on the request message timeout is defined per session. In other words, a request message timeout for one session cannot cause the failover for some other session.

  The maximum number of transmissions per session is hard-coded to 2 and the same message is never re-transmitted to the same TCP socket (a TCP socket is defined as a current peering connection defined by the TCP source/destination IP addresses/ports; closing and then reopening a connection to the same peer will result in creation of a new TCP socket). Once the original message for the session times out on the primary peer, the message is re-transmitted to the secondary peer, provided that the secondary peer is available and the failover is enabled with the corresponding handling mechanism. In case that the secondary peer is unavailable, the original message is not be re-transmitted to the same primary peer again.

  Once the reply from a peer is received, the session is tied to that peer until the next timeout. In other words, the session always sticks to the peer from which it received the last response.

- The handling behavior if the response from the peer is not received or the peers are not available at all (all peering connections are closed). If the response to a session initiation message is not received, the fate of the session will depend on the configuration (the session can be terminated or continue to exist with default parameters).

Default on-failure failover enabled handling terminate

Parameters **failover enabled** — Specifies that application session is allowed to switch to the secondary peer

failover disabled — Specifies that the application session is not allowed to switch to the secondary peer

handling continue — Specifies that the sessions continue to exist if the response to a transmitted CCR message is not received. Whether the transmitted message is re-transmitted depends on the failover configuration. In case of session initiation procedure in the Gx case (CCR-I timeout), the subscriber host is nonetheless instantiated with the default parameters, assuming that they are provided. If the default parameter are not provided, the subscriber host initiation fails.

handling retry-and-terminate — Specifies that the message is re-transmitted if the peer-failover is enabled and the secondary peer is available. Once the retransmitted request message is timed out, the application session is terminated.

handling terminate — Specifies that the session is terminated if the response to the original request message is not received. No re-transmissions is attempted, regardless if the failover is enabled.

tx-timer

Syntax **tx-timer** *seconds*
no tx-timer

Context config>subscr-mgmt>diam-app-plcy

Description This command defines the time-out period for the application's request messages (CCR-I/U/T).

The **on-failure** configuration determines the action that will be taken once the message times out.

The **no** form of this command reverts to the default value.

Default tx-timer 10

Parameters *seconds* — Specifies the Tx Timer value (in seconds) for this policy.

Values 10 to 1000

diameter-application-policy

Syntax **diameter-application-policy** *policy-name*
no diameter-application-policy

Context config>service>vpls>sap
config>service>vprn>sub-if>grp-if
config>service>ies>sub-if>grp-if
config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>pppoe>host

Description This command applies the **diameter-application-policy** to the processing of the host attachment requests.

The **no** form of this command reverts to the default value.

Parameters *policy-name* — Specifies the name of the diameter policy, up to 32 characters

9.25.2.1.8 Diameter Base Commands

diameter

Syntax **diameter**

Context config>aaa

Description This command enables the context to configure Diameter parameters.

node

Syntax **node** *origin-host-string* [**origin-realm** *origin-realm-string*] [**create**]
no node *origin-host-string*

Context config>aaa>diam>node

Description This command creates a Diameter client node in the SR OS. Multiple Diameter client nodes with their own peer definitions are simultaneously supported in SR OS.

Each such node is defined by a unique DiameterIdentity (the origin host and realm names).

The **no** form of the command removes the origin host string from the configuration.

Parameters *origin-host-string* — The origin host name, up to 80 characters, is a mandatory parameter that translates to an Origin-Host AVP that is carried in all Diameter messages. The origin host and origin realm form a Diameter Identity that must be unique within the Diameter network in which they participate.

origin-realm-string — The origin realm name, up to 80 characters, is an optional parameter that translates to an Origin-Realm AVP that is carried in all Diameter messages. The origin host and origin realm form a Diameter Identity that must be unique within the Diameter network in which they participate

If the origin realm name is not configured, it will be extracted from the origin host parameter.

create — Keyword used to create the Diameter client node. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

connection-timer

Syntax **connection-timer** *seconds*
no connection-timer

Context configure>aaa>diam>node
config>aaa>diam>node>peer

Description This command configures the Diameter node connection timer that defines the time the systems waits before attempting to reconnect to a peer after the connection was lost.

The **no** form of the command reverts to the default.

Default connection-timer 30

Parameters *seconds* — Specifies the Diameter node connection timer.

Values 1 to 1000

peer

Syntax **peer index** *index* [*destination-host-string*] [**create**]
no peer index *index*

Context configure>aaa>diam>node

Description This command creates context for diameter peer configuration within a Diameter client node in SR OS. Up to five Diameter peers can be configured within a given Diameter client node.

This command is not applicable to legacy Diameter base.

The **no** form of the command removes the peer index information from the configuration.

Parameters *index* — Specifies the index of a peer. Index is used to break a tie if a Diameter route for a given host or realm destination points to multiple diameter peers with the same preference. In such scenario, the peer with the lowest index will be selected as next-hop in traffic forwarding.

Values 1 to 5

destination-host-string — This mandatory parameter identifies the peer by its name or realm (of type DiameterIdentity). This peer name must much the one in Origin-Host AVP and Origin-Realm AVP received in Capability Exchange Answer message. In case of a mismatch, the TCP connection will be terminated.

create — Keyword used to create the peer index. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

address

Syntax **address** [*ip-address* | *ipv6-address*]
no address

Context config>aaa>diam>node>peer

Description This command configures IPv4 or IPv6 address for a Diameter peer.

The **no** from of the command removes the IPv4 or IPv6 from the peer configuration.

Parameters *ip-address* — Specifies the IPv4 address in the a.b.c.d form

ipv6-address — Specifies the IPv6 address in the form:

x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
where:
x - [0 to FF]
d - [0 to 255]

connection-timer

Syntax [**no**] **connection-timer** *connection-time*

Context config>aaa>diam-peer-plcy
config>aaa>diam-peer-plcy>peer

Description This command defines the frequency of attempts to open a TCP connection to each peer within a given Diameter client node. Once a TCP connection fails to be established (due to timeouts of TCP SYN packets) or an existing TCP connection fails, the next attempt to open the connection will be tried upon the expiry of the connection-timer. There is no limit on the number of attempts.

The **no** form of this command reverts to the default.

Default connection-timer 30

Parameters *connection-time* — Specifies the amount of time, in seconds, for the diameter peer policy connection

Values 1 to 1000

preference

Syntax **preference** *preference*
no preference

Context config>aaa>diam>node>peer

Description This command configures the Diameter routing preference for a peer. All open peers are installed in the Diameter realm routing table but only one with the highest preference will be used as the next hop for a given destination realm. If multiple peers with the same preference are configured for the same realm, the peer index will be used to break the tie.

The **no** form of the command removes the preference from the configuration.

Parameters *preference* — Specifies the peer preference.

Values 1 to 100

watchdog-timer

Syntax **watchdog-timer** *seconds*
no watchdog-timer

Context config>aaa>diam>node>peer

Description Watchdog messages are used to verify liveliness of a Diameter peer. A single watchdog messages will be sent to a peer in case that no traffic is received from it within a configured watchdog-timer. This watchdog request message will solicit a watchdog answer message from the peer. If no traffic (watchdog answer or otherwise) is received from the peer in response to watchdog request while the watchdog timer is running, the peer will be put in suspicious state and a peer failover routine will be triggered.

The peer will be closed after it has been in suspicious mode for the duration of one more watchdog-timer interval without receiving any traffic from it.

Once the peer is recovered, it will take 3 successful exchanges of diameter watchdog messages (DWR/DWA) for the peer to becomes used again in Diameter forwarding. This behavior is described in RFC 3539, §3.4.1, *Authentication, Authorization and Accounting (AAA) Transport Profile*).

This command is not applicable to legacy implementations of Diameter base in the SR OS.

The **no** form of the command removes the watchdog timer value from the configuration.

Default watchdog-timer 30

Parameters *seconds* — Specifies the device watchdog timer, in seconds, used on all connections by this peer

Values 1 to 1000

router

Syntax **router** *router-instance*
router service *service-name*
no router

Context config>aaa>diam-peer-plcy
config>aaa>diam>node

Description This command references the routing-instance from which diameter peering connection is initiated.

The **no** form of the command reverts to the default

Parameters *router-instance* — specifies the router instance

Values *router-name* | *vprn-svc-id*
router-name: Base, management Default - Base
vprn-svc-id: 1 to 2147483647

service-name — Specifies the service name, up to 64 characters

diameter-peer-policy

Syntax **diameter-peer-policy** *peer-policy-name* [**role** {**client** | **proxy**}] [**create**]
no diameter-peer-policy *peer-policy-name*

Context config>aaa

Description This command is the foundation of the legacy Diameter base implementation in the SR OS node. It defines a Diameter client node in the SR OS and within, where the peering connections are defined. Various diameter applications (NASREQ, Gx, and Gy) can reference this policy.

This command applies only to the legacy Diameter base.

The **no** form of the command removes the Diameter node origin host string and destination realm from the configuration.

Parameters *peer-policy-name* — Specifies the peer policy name, up to 32 characters

role client — Specifies that Diameter is configured as a client. The client initiates peering connections towards the server or Diameter proxy. Various applications such as NASREQ, Gx, and Gy are layered directly on top of the Diameter client.

Default role client

role proxy — Specifies that the SR OS node is configured as an intermediary Diameter node in dual-homed topology. In such topology, there are two Diameter proxy nodes placed between dual-homed Diameter clients on one side and Diameter server on the other. Diameter proxy can assume an active or standby state.

create — Keyword used to create the Diameter peer policy. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

diameter-peer-policy

Syntax **diameter-peer-policy** *peer-policy-name*
no diameter-peer-policy

Context configure>subscr-mgmt>diam-app-pol

Description The **diameter-peer-policy** command contain a legacy Diameter client node definition in SR OS is referenced by an application (NASREQ, Gx, and Gy). The application will then use this **diameter-peer-policy** to communicate with external Diameter nodes.

This command applies only to legacy implementations of Diameter base and it is mutually exclusive with the **diameter-node** command under the same hierarchy.

The **no** form of the command removes the peer policy name from the configuration.

Parameters *peer-policy-name* — Specifies the peer policy name, up to 32 characters

diameter-proxy

Syntax [**no**] **diameter-proxy**

Context configure>redundancy>multi-chassis>peer>sync

Description This command enables Diameter redundancy in dual-homed environment. This command is used to synchronize Diameter proxy parameters between two redundant SR OS nodes. The parameters that are synchronized are:

- chassis system MAC
- Diameter proxy states
- Origin-State-Id

This command is applicable only to legacy implementations of Diameter base in the SR OS.

The **no** form of the command disables Diameter redundancy.

9.25.2.1.9 Explicit Subscriber Mapping Commands

explicit-subscriber-map

Syntax **explicit-subscriber-map**

Context config>subscriber-mgmt

Description This command configures an explicit subscriber mapping

entry

Syntax **entry key** *sub-ident-string* [**sub-profile** *sub-profile-name*] [**alias** *sub-alias-strin***g**] [**sla-profile** *sla-profile-name*] [**app-profile** *app-profile-name*]
no entry key *sub-profile-string*

Context config>subscr-mgmt>explicit-sub-map

Description This command configures a subscriber identification string.

The **no** form of this command reverts to the default.

Parameters *sub-ident-string* — Specifies the profile string, up to 32 characters.

sub-profile-name — Specifies an existing subscriber profile name, up to 32 characters.

sub-alias-string — Specifies an alias for the subscriber identification string, up to 64 characters.

sla-profile-name — Specifies an existing SLA profile, up to 32 characters.

app-profile-name — Specifies an app profile name up to 256 characters.

9.25.2.1.10 GSMP Configuration Commands

Note: The commands in this section apply only to the 7750 SR.

gsmp

Syntax **gsmp**

Context config>service>vpls
config>service>vprn

Description This command enables the context to configure GSMP connections maintained in this service.

group

Syntax **group** *name* [**create**]
no group *name*

Context config>service>vpls>gsmp
config>service>vprn>gsmp

Description This command specifies a GSMP name. A GSMP group name is unique only within the scope of the service in which it is defined.

The **no** form of this command reverts to the default.

Parameters *name* — Specifies a GSMP name up to 32 characters.

create — Keyword used to create the GSMP group name. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

ancp

Syntax **ancp**

Context config>service>vpls>gsmp>group
config>service>vprn>gsmp>group

Description This command enables the context to configure ANCP parameters for this GSMP group.

dynamic-topology-discover

Syntax [**no**] **dynamic-topology-discover**

Context config>service>vpls>gsmp>group>ancp
config>service>vprn>gsmp>group>ancp

Description This command enables the ANCP dynamic topology discovery capability.

The **no** form of this command disables the feature.

oam

Syntax [**no**] **oam**

Context config>service>vpls>gsmp>group>ancp

config>service>vprn>gsmp>group>ancp

Description This command enables the GSMP ANCP OAM capability to be negotiated at startup of the GSMP connection.

The **no** form of this command disables the feature.

hold-multiplier

Syntax **hold-multiplier** *multiplier*
no hold-multiplier

Context config>service>vpls>gsmp
config>service>vprn>gsmp

Description This command configures the hold-multiplier for the GSMP connections in this group.

The **no** form of this command removes the multiplier value from the GSMP configuration.

Parameters *multiplier* — Specifies the GSMP hold multiplier value.

Values 1 to 100

idle-filter

Syntax [**no**] **idle-filter**

Context config>service>vpls>gsmp>group
config>service>vprn>gsmp>group

Description This command when applied will filter out new incoming ANCP messages while the subscriber DSL-line-state is idle. The command takes effect at the time that it is applied. Existing subscribers already in idle state are not purged from the database.

The **no** form of this command reverts to the default.

keepalive

Syntax **keepalive** *seconds*
no keepalive

Context config>service>vpls>gsmp>group
config>service>vprn>gsmp>group

Description This command configures keepalive values for the GSMP connections in this group.

The **no** form of this command reverts to the default.

Parameters *seconds* — Specifies the GSMP keepalive timer value in seconds.

Values 1 to 25

neighbor

Syntax **neighbor** *ip-address* [**create**]
no neighbor *ip-address*

Context config>service>vpls>gsmp>group
config>service>vprn>gsmp>group

Description This command enables the context to configure a GSMP ANCP neighbor parameters.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the IP address of the GSMP ANCP neighbor.

create — Keyword used to create the neighbor instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

local-address

Syntax **local-address** *ip-address*
no local-address

Context config>service>vpls>gsmp>group>neighbor
config>service>vprn>gsmp>group>neighbor

Description This command configures the source ip-address used in the connection towards the neighbor. The local address is optional. If specified the node will accept connections only for that address in the service running ANCP. The address may be created after the reference but connections will not be accepted until it is created. If the local address is not used, the system accepts connections on any interface within the routing context.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the source IP address to be used in the connection toward the neighbor.

priority-marking

Syntax **priority-marking dscp** *dscp-name*
priority-marking prec *ip-prec-value*
no priority-marking

Context config>service>vpls>gsmp>group>neighbor

config>service>vprn>gsmp>group>neighbor

Description This command configures the type of priority marking to be used.

The **no** form of this command reverts to the default.

Parameters *dscp-name* — Specifies the DSCP code-point to be used.

Values be, cp1, cp2, cp3, cp4, cp5, cp6, cp7, cs1, cp9, af11, cp11, af12, cp13, af13, cp15, cs2, cp17, af21, cp19, af22, cp21, af23, cp23, cs3, cp25, af31, cp27, af32, cp29, af33, cp31, cs4, cp33, af41, cp35, af42, cp37, af43, cp39, cs5, cp41, cp42, cp43, cp44, cp45, ef, cp47, nc1, cp49, cp50, cp51, cp52, cp53, cp54, cp55, nc2, cp57, cp58, cp59, cp60, cp61, cp62, cp63

ip-prec-value — Specifies the precedence value to be used.

Values 0 to 7

persistency-database

Syntax [**no**] **persistency-**database

Context config>service>vpls>gsmp>group
config>service>vprn>gsmp>group

Description This command enables the system to store DSL line information in memory. If the GSMP connection terminates, the DSL line information will remain in memory and accessible for RADIUS authentication and accounting.

The **no** form of this command reverts to the default.

9.25.2.1.11 Host Lockout Commands

Note: The commands described in this section apply only to the 7750 SR.

host-key

Syntax **host-key** {**mac**}
no host-key

Context config>subscr-mgmt>host-lockout-plcy

Description This command specifies the parameters used in host identification for lockout on a given SAP or capture SAP.

no host-key – include (MAC address, Circuit-Id, Remote-Id)

host-key mac – include MAC address only

"host-key mac" should be used in DHCPv4 scenarios where Circuit-Id and Remote-Id are changed with "dhcp option action replace" configuration: a host lockout context is created with the replaced Circuit-Id/Remote-Id; with the default host-key (including Circuit-Id and Remote-Id), lockout does not kick in on the original trigger packet when it is retransmitted by the client.

Changing the host-key to mac should be used with care: all hosts with the same MAC address on a given SAP or capture SAP are identified as a single host with respect to host-lockout.

This command cannot be changed when the **host-lockout-policy** is referenced (configured under a SAP context).

The **no** form of this command reverts to the default value.

Parameters **mac** — Specifies to use the MAC address only for host identification for lockout.

host-lockout-policy

Syntax **host-lockout-policy** *policy-name* [**create**]
no host-lockout-policy *policy-name*

Context config>subscr-mgmt

Description This command creates a host lockout policy. The policy contains set of host lockout configuration parameters. It is applied to SAP or MSAPs (by a MSAP-policy). Any change does not impact existing locked-out hosts, but only new incoming hosts that enter lockout.

The **no** form of this command removes the policy name from the configuration. The policy must not be associated with any entity.

Parameters *policy-name* — Specifies an existing host lockout policy to associate with the SAP.

create — Specifies the keyword required to create the host lockout policy. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

address

Syntax **address** {*ip-address/mask* | *ip-address netmask*} [**broadcast all-ones** | **host-ones**] [**track-srrp** *srrp-instance*]
no address [*ip-address/mask* | *ip-address netmask*]

Context config>service>ies>if
config>service>vprn>if

Description This command assigns an IP address, IP subnet, and broadcast address format to an IES IP router interface. Only one IP address can be associated with an IP interface.

An IP address must be assigned to each IES IP interface. An IP address and a mask are used together to create a local IP prefix. The defined IP prefix must be unique within the context of the routing instance. It cannot overlap with other existing IP prefixes defined as local subnets on other IP interfaces in the same routing context within the router.

The local subnet that the **address** command defines must be part of the services' address space within the routing context using the **config router service-prefix** command. The default is to disallow the complete address space to services. Once a portion of the address space is allocated as a service prefix, that portion can be made unavailable for IP interfaces defined within the **config router interface** CLI context for network core connectivity with the **exclude** option in the **config router service-prefix** command.

The IP address for the interface can be entered in either CIDR (Classless Inter-Domain Routing) or traditional dotted decimal notation. The show commands display CIDR notation and is stored in configuration files.

By default, no IP address or subnet association exists on an IP interface until it is explicitly created.

Table 108 **Address Admin and Operational States**

Address	Admin State	Oper State
No address	up	down
No address	down	down
1.1.1.1	up	up
1.1.1.1	down	down

The operational state is a read-only variable and the only controlling variables are the address and admin states. The address and admin states are independent and can be set independently. If an interface is in an administratively up state and an address is assigned, it becomes operationally up and the protocol interfaces and the MPLS LSPs associated with that IP interface are reinitialized.

The **no** form of this command removes the IP address assignment from the IP interface. When the **no address** command is entered, the interface becomes operationally down.

Parameters *ip-address* — Specifies the IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that is used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

/ — Specifies a parameter delimiter and separates the *ip-address* portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the "*/*" and the *mask-length* parameter. If a forward slash is not immediately following the *ip-address*, a dotted decimal mask must follow the prefix.

mask-length — Specifies the subnet mask length when the IP prefix is specified in CIDR notation. When the IP prefix is specified in CIDR notation, a forward slash (/) separates the *ip-address* from the *mask-length* parameter. The mask length parameter indicates the number of bits used for the network portion of the IP address; the remainder of the IP address is used to determine the host portion of the IP address.

Note: A mask length of 32 is reserved for loopback addresses (includes system addresses).

Default 0 to 31

mask — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that is used in a logical 'AND' function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252.

Note: A mask of 255.255.255.255 is reserved for system IP addresses.

broadcast — Overrides the default broadcast address used by the IP interface when sourcing IP broadcasts on the IP interface. If no broadcast format is specified for the IP address, the default value is **host-ones** which indicates a subnet broadcast address. Use this parameter to change the broadcast address to **all-ones** or revert back to a broadcast address of **host-ones**.

The broadcast format on an IP interface can be specified when the IP address is assigned or changed.

This parameter does not affect the type of broadcasts that can be received by the IP interface. A host sending either the local broadcast (**all-ones**) or the valid subnet broadcast address (**host-ones**) is received by the IP interface.

Default host-ones

all-ones — Specifies the broadcast address used by the IP interface for this IP address is 255.255.255.255, also known as the local broadcast.

host-ones — Specifies that the broadcast address used by the IP interface for this IP address is the subnet broadcast address. This is an IP address that corresponds to the local subnet described by the *ip-address* and the *mask-length* or *mask* with all the host bits set to binary one. This is the default broadcast address used by an IP interface.

The **broadcast** parameter within the **address** command does not have a negate feature, which is usually used to revert a parameter to the default value. To change the **broadcast** type to **host-ones** after being changed to **all-ones**, the **address** command must be executed with the **broadcast** parameter defined.

allow-directed-broadcasts

Syntax [**no**] **allow-directed-broadcasts**

Context config>router>if
config>service>ies>if
config>service>vprn>if

Description This command enables the forwarding of directed broadcasts out of the IP interface.

A directed broadcast is a packet received on a local router interface destined for the subnet broadcast address on another IP interface. The **allow-directed-broadcasts** command on an IP interface enables or disables the transmission of packets destined to the subnet broadcast address of the egress IP interface.

When enabled, a frame destined to the local subnet on this IP interface is sent as a subnet broadcast out this interface. Care should be exercised when allowing directed broadcasts as it is a well-known mechanism used for denial-of-service attacks.

When disabled, directed broadcast packets discarded at this egress IP interface are counted in the normal discard counters for the egress SAP.

By default, directed broadcasts are not allowed and are discarded at this egress IP interface.

The **no** form of this command disables the forwarding of directed broadcasts out of the IP interface. All broadcasts are dropped.

sap

Syntax [**no**] **sap** *sap-id*

Context config>service>ies
config>service>vprn
config>service>ies>if
config>service>vprn>if
config>service>ies>sub-if
config>service>vprn>sub-if

Description This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the router. Each SAP must be unique.

The **no** form of this command removes the SAP ID from the configuration.

For further information about the **sap** command usage, syntax and description, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* and the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN*.

cflowd-parameters

Syntax **cflowd-parameters**

Context config>service>ies>if

Description This command enables the context to configure cflowd parameters on the interface.

sampling

Syntax **sampling** {**unicast** | **multicast**} **type** {**acl** | **interface**} [**direction** {**ingress-only** | **egress-only** | **both**}]
no sampling {**unicast** | **multicast**}

Context config>service>ies>if>cflowd
config>service>vprn>if>cflowd

Description This command configures cflowd sampling parameters.

The **no** form of this command removes the sampling parameters.

Parameters **unicast | multicast** — Specifies unicast or multicast sampling.

type — Specifies the cflowd sampling type on the given virtual router interface.

Values acl — Specifies ACL cflowd analysis be applied to the given virtual router interface.

interface — Specifies interface cflowd analysis be applied to the given virtual router interface

direction — Specifies the direction of the cflowd analysis that is applied to the given virtual router interface.

Values ingress-only — Specifies an ingress only direction of the cflowd analysis be applied to the given virtual router interface.

egress-only — Specifies an egress only direction of the cflowd analysis be applied to the given virtual router interface.

both — Specifies both ingress and egress direction of the cflowd analysis be applied to the given virtual router interface.

host-lockout-policy

Syntax **host-lockout-policy** *policy-name*
no host-lockout-policy

Context config>service>ies>if>sap
config>service>ies>sub-if>sap
config>service>vpls>sap
config>service>vprn>if>sap
config>service>vprn>sub-if>sap

Description This command selects an existing host lockout policy. The **host-lockout-policy** *policy-name* is created in the **config>subscr-mgmt** context.

The **no** form of this command removes the policy name from the SAP configuration.

Parameters *policy-name* — Specifies an existing host lockout policy, up to 32 characters, to associate with the SAP.

lockout-time

Syntax **lockout-time** [**min** *seconds*] [**max** *seconds*]
no lockout-time

Context config>subscr-mgmt>host-lockout-plcy

Description This command configures the time for which a client stays in the lockout state during which authentication and ESM host creation is suppressed.

The **no** form of this command reverts to the default value.

Default lockout-time min 10 max 3600

Parameters **min** *seconds* — Specifies the minimum lockout-time for this host lockout policy.

Values 1 to 86400

Default 10 seconds

max *seconds* — Specifies the maximum lockout-time for this host lockout policy.

Values 1 to 86400

Default 3600 seconds

lockout-reset-time

Syntax **lockout-reset-time** *seconds*
no lockout-reset-time

Context config>subscr-mgmt>host-lockout-plcy

Description This command configures the time that needs to elapse from the point a client enters lockout to when the client's lockout time can be reset to the configured minimum value. The range is 1 second.

The **no** form of this command reverts to the default value.

Default lockout-reset-time 60

Parameters *seconds* — Specifies the lockout reset time, in seconds.

Values 1 to 86400

max-lockout-hosts

Syntax **max-lockout-hosts** *hosts*
no max-lockout-hosts

Context config>subscr-mgmt>host-lockout-plcy

Description When a client enters lockout, authentication and ESM host creation is suppressed. A lightweight context maintains the lockout state and the timeouts for the client in lockout. This command allows the number of lockout contexts to be configured per SAP. If the number of existing contexts reaches the configured count, incoming hosts that fail authentication or creation are not subject to lockout, and are retired as normal.

The **no** form of this command reverts to the default value.

Default max-lockout-hosts 100

Parameters *hosts* — Specifies the maximum number of lockout host.

Values 1 to 32000

firewall-policy

Syntax **firewall-policy** *policy-name*
no firewall-policy

Context config>subscr-mgmt>sub-prof

Description This command specifies the firewall policy for subscribers associated with this subscriber profile.

The **no** form of this command removes the firewall policy name.

Parameters *policy-name* — Specifies the name of an existing firewall policy up to 32 characters.

host-tracking-policy

Syntax **host-tracking-policy** *policy-name* [**create**]
no host-tracking-policy *policy-name*

Context config>subscr-mgmt
config>subscr-mgmt>sub-prof

Description This command configures a host tracking policy. IGMP host tracking is an option in the subscriber profile that allows the factoring in of a subscriber's (multicast) video traffic by reducing the unicast operational egress aggregate rate or the rate of the scheduler specified in the ANCP policy to account for a subscriber's multicast traffic. If no ANCP policy is defined, the egress aggregate rate configured in the subscriber profile is reduced. If an ANCP policy is defined, the **rate-modify** command in the policy specifies whether the egress aggregate rate or the rate of the egress policer specified in the policy is to be reduced to account for the subscriber's multicast traffic.

The **no** form of this command reverts to the default value.

Parameters *policy-name* — Specifies a host tracking policy name up to 32 characters.

egress-rate-modify

Syntax **egress-rate-modify agg-rate-limit**
egress-rate-modify scheduler *scheduler-name*
no egress-rate-modify

Context config>subscr-mgmt>trk-plcy

Description This command specifies the egress-rate modification that is to be applied.

The **no** form of this command reverts to the default value.

Parameters **agg-rate-limit** — Specifies to use the egress rate limit.

scheduler-name — Specifies the scheduler name up to 32 characters.

9.25.2.1.12 IGMP Policy Commands

Note: The commands described in this section apply only to the 7750 SR.

igmp-policy

Syntax **igmp-policy** *policy-name* [**create**]
no igmp-policy

Context config>sub-mgmt

Description This command configures an IGMP policy.

The **no** form of this command reverts to the default value.

Parameters *policy-name* — Specifies the policy name up to 32 characters.

create — Keyword used to create the IGMP policy. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

disable-router-alert-check

Syntax [**no**] **disable-router-alert-check**

Context config>sub-mgmt>igmp-policy

Description This command disables router alert checking for IGMP messages received on this interface.

The **no** form of this command reverts to the default value.

egress-rate-modify

Syntax **egress-rate-modify** [**agg-rate-limit** | **scheduler** *scheduler-name*]
no egress-rate-modify

Context config>sub-mgmt>igmp-policy

Description This command is used to apply HQoS Adjustment to a subscriber. HQoS Adjustment is needed when multicast traffic flow for the subscriber is dissociated from subscriber host queues. Multicast redirection is typical such case although it can be applied in direct IPoE subscriber per-sap replication mode.

The channel bandwidth definition policy is defined in the mcac policy under the **config>router>mcac>policy** context. The policy is applied under the redirected interface or under the group-interface.

In order for HQoS Adjustment to take effect, sub-mcac-policy must be in a no shutdown mode and applied under the sub-profile even if mcac is not deployed.

The **no** form of this command reverts to the default value.

Parameters **agg-rate-limit** — Specifies the aggregate14 rate modification to be applied. The subscriber's bandwidth is capped via the **agg-rate-limit** command in the sub-profile or with a Change of Authorization (CoA) request. This bandwidth cap is dynamically adjusted according to the multicast channel definition and channel association with the host via IGMP.

scheduler-name — Specifies the schedule name. The subscriber's bandwidth is capped via the scheduling-policy in the sub-profile or with a Change of Authorization (CoA) request. HQoS Adjustment will modify the rate of the scheduler defined in the scheduling policy or configured via CoA.

import

Syntax **import** *policy-name*
no import

Context config>sub-mgmt>igmp-policy

Description This command specifies the import policy to filter IGMP packets.

The **no** form of this command reverts to the default value.

Parameters *policy-name* — Specifies the policy name up to 32 characters.

max-num-groups

Syntax **max-num-groups** *max-num-groups*
no max-num-groups

Context config>sub-mgmt>igmp-policy

Description This command configures the maximum number of multicast groups.

The **no** form of this command reverts to the default value.

Parameters *max-num-groups* — Specifies the maximum number of multicast groups.

Values 0 to 16000

max-num-sources

Syntax **max-num-sources** *max-num-sources*
no max-num-sources

Context config>sub-mgmt>igmp-policy

Description This command configures the maximum number of multicast sources.

The **no** form of this command disables the command.

Parameters *max-num-sources* — Specifies the maximum number of multicast sources.

Values 1 to 1000

max-num-grp-sources

Syntax **max-num-grp-sources** [*max-num-grp-sources*]
no max-num-grp-sources

Context config>sub-mgmt>igmp-policy
config>sub-mgmt>msap-policy>igmp-host-tracking

Description This command configures the maximum number of group sources for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of group sources, the group sources that are already accepted are not deleted. Only new group sources will not be allowed. When this object has a value of 0, there is no limit to the number of group sources.

The **no** form of this command removes the value from the configuration.

Parameters *max-num-grp-sources* — Specifies the maximum number of multicast sources allowed to be tracked per group.

Values 1 to 32000

mcast-reporting

Syntax [**no**] **mcast-reporting**

Context config>sub-mgmt>igmp-policy

Description This command enables the context to configure mcast reporting.

mcast-reporting-dest

Syntax **mcast-reporting-dest** *dest-name*
no mcast-reporting-dest

Context config>subscr-mgmt>igmp-policy>mcast-reporting
config>subscr-mgmt>host-tracking-policy>mcast-reporting

Description This command references the multicast reporting destination to which IGMP related events are exported.

The multicast reporting destination is referenced with the subscriber itself or within the Host-Tracking-Policy.

The **no** form of this command reverts to the default value.

Parameters *dest-name* — Specifies the name of the multicast reporting destination up to 32 characters.

opt-reporting-fields

Syntax **opt-reporting-fields** [**host-mac**] [**pppoe-session-id**] [**svc-id**] [**sap-id**]
no opt-reporting-fields

Context config>subscr-mgmt>igmp-policy>mcast-reporting
config>subscr-mgmt>host-tracking-policy>mcast-reporting

Description This command specifies optional data relevant to the IGMP event that can be exported. This optional data includes:

- Host MAC address
- PPPoE session-ID
- Service ID
- SAP

The **no** form of this command reverts to the default value.

Parameters **host-mac** — Specifies that the host-mac optional field should be included into the multicast reporting messages.

pppoe-session-id — Specifies that the pppoe-session-id optional field should be included into the multicast reporting messages.

svc-id — Specifies that the svc-id optional field should be included into the multicast reporting messages.

sap-id — Specifies that the sap-id optional field should be included into the multicast reporting messages.

query-interval

Syntax **query-interval** *seconds*
no query-interval

Context config>subscr-mgmt>igmp-policy

Description This command specifies the frequency at which the querier router transmits general host-query messages. Host-query messages solicit group membership information and are sent to the all-systems multicast group address, 224.0.0.1. If nothing is configured, by default, the **query-interval** takes the value defined in the **config**>**router**>**igmp** context or in the **config**>**service**>**vprn**>**igmp** context. It is highly recommended that all three query intervals be configured together on each IGMP policy.

The **no** form of this command reverts to the default value.

Parameters *seconds* — Specifies the frequency, in seconds, at which the router transmits general host-query messages.

Values 2 to 1024

query-last-member-interval

Syntax **query-last-member-interval** *seconds*
no query-last-member-interval

Context config>subscr-mgmt>igmp-policy

Description This command configures the frequency at which the querier router sends group-specific query messages, including the messages sent in response to leave-group messages. The shorter the interval, the faster a loss of the last member of a group can be detected. If nothing is configured, by default, the **query-last-member-interval** takes the value defined in the **config**>**router**>**igmp** context or the **config**>**service**>**vprn**>**igmp** context. It is highly recommended that all three query intervals be configured together on each IGMP policy.

The **no** form of this command reverts to the default value.

Parameters *seconds* — Specifies the frequency, in seconds, at which query messages are sent.

Values 1 to 1024

query-response-interval

Syntax **query-response-interval** *seconds*
no query-response-interval

Context config>subscr-mgmt>igmp-policy

Description This command configures the query response interval. If nothing is configured, by default, the **query-response-interval** takes the value defined in the **config**>**router**>**igmp** context or in the **config**>**service**>**vprn**>**igmp** context. It is highly recommended that all three query intervals be configured together on each IGMP policy.

The **no** form of this command reverts to the default value.

Parameters *seconds* — Specifies the length of time, in seconds, that the querier router waits to receive a response to from the host.

Values 1 to 1023

sub-mcac-policy

Syntax **sub-mcac-policy** *policy-name*
no sub-mcac-policy

Context config>subscr-mgmt>sub-profile

Description This command enables the context to configure sub MCAC policy parameters. The policy template in which the MCAC bandwidth limits are defined. MCAC for the subscriber is effectively enabled with this command when the sub-profile is applied to the subscriber. The bandwidth of the channels is defined in a different policy (under the **config>router>mcac** context) and this policy is applied on the interface level as follows:

- for regular interfaces in the **config>service/router>igmp>interface>mcac** context
- for HQoS adjustment, it is mandatory that the sub-mcac-policy be created and applied to the subscriber. The sub-mac-policy does not have to contain any bandwidth constrains, but it has to be in a no shutdown state in order for HQoS Adjustment to work.

Parameters *policy-name* — Specifies the policy name up to 32 characters.

version

Syntax **version** *version*
no version

Context config>subscr-mgmt>igmp-policy

Description This command configures the version of IGMP.

The **no** form of this command reverts to the default value.

Default version 3

Parameters *version* — Specifies the IGMP version.

Values 1, 2 or 3

fast-leave

Syntax [**no**] **fast-leave**

Context config>sub-mgmt>igmp-policy

Description This command enables IGMP fast-leave processing.

The **no** form of this command reverts to the default value.

Default fast-leave

static

Syntax **static**

Context config>sub-mgmt>igmp-policy

Description This command enables the context to configure IGMP static group membership parameters.

per-host-replication

Syntax **per-host-replication** [**uni-mac** | **mcast-mac**]
no per-host-replication

Context config>subscr-mgmt>igmp-policy

Description This command enables per-host-replication in IPoE model. For PPPoX, per-host-replication is the only mode of operation. In the per-host-replication mode, multicast traffic is replicated per each host within the subscriber irrespective of the fact that some hosts may be subscribed to the same multicast stream. As a result, in case that multiple hosts within the subscriber are registered for the same multicast group, the multicast streams of that group is generated. The destination MAC address of multicast streams is changed to unicast so that each host receives its own copy of the stream. Multicast traffic in the per-host-replication mode can be classified via the existing QoS CLI structure. As such the multicast traffic will flow through the subscriber queues. HQoS Adjustment is not needed in this case.

The alternative behavior for multicast replication in IPoE environment is per-SAP- replication. In this model, only a single copy of the multicast stream is sent per SAP, irrespective of the number of hosts that are subscribed to the same multicast group. This behavior applies to 1:1 connectivity model as well as on 1:N connectivity model (SAP centric behavior as opposed to subscriber centric behavior).

In the per-SAP-replication model the destination MAC address is multicast (as opposed to unicast in the per-host-replication model). Multicast traffic is flowing via the SAP queue which is outside of the subscriber context. The consequence is that multicast traffic is not accounted in the subscriber HQoS. In addition, HQoS Adaptation is not supported in the per SAP replication model.

Default no per-host-replication — By default there is no per host replication and replication is done per SAP. This mode utilizes the SAP queues. With per-host-replication it will allow the use of the subscriber queues. Per-host-replication uses unicast MAC and multicast IP to deliver multicast content to end hosts. This is useful for multi host per SAP cases. To interoperate with end devices that do not support unicast MAC, there is an option to use per-host-replication with a multicast MAC. The traffic is the same as replication per SAP but the difference of using the subscriber queues.

Parameters **uni-mac** — Specifies that multicast traffic is sent with a unicast MAC and multicast IP.

mcast-mac — Specifies that multicast traffic is sent with a multicast MAC and IP.

redirection-policy

Syntax **redirection-policy** *policy-name*
no redirection-policy

Context config>sub-mgmt>igmp-policy

Description This command will apply multicast redirection action to the subscriber. The redirection action along with the redirected interface (and possibly service id) is defined in the referenced policy-name. IGMP messages is redirected to an alternate interface if that alternate interface has IGMP enabled. The alternate interface does not have to have any multicast groups registered via IGMP. Currently all IGMP messages are redirected and there is no ability to selectively redirect IGMP messages based on match conditions (multicast-group address, source IP address, and so on). Multicast redirection is supported between VPRN services and also between interfaces within the Global Routing Context. Multicast Redirection is not supported between the VPRN services and the Global Routing Table (GRT).

IGMP state is maintained per subscriber host and per redirected interface. Traffic is however forwarded only on the redirected interface.

The **no** form of this command reverts to the default value.

Parameters *policy-name* — Specifies the redirection policy to be applied to this host IGMP policy up to 32 characters. This is a policy defined in the **config>router>policy-option>policy-statement** context.

group

Syntax [**no**] **group** *ip-address*

Context config>sub-mgmt>igmp-policy>static

Description This command adds or removes a static multicast group.

The **no** form of this command reverts to the default value.

Parameters *ip-address* — Specifies the multicast group IP address.

Values a.b.c.d

9.25.2.1.13 Managed SAP Policy Commands

msap-policy

Syntax **msap-policy** *msap-policy-name* [**create**]
no msap-policy *msap-policy-name*

Context config>subscr-mgmt

Description This command configures a managed SAP policy. Managed SAPs allow the use of policies and a SAP template for the creation of a SAP.

The **no** form of this command removes the MSAP policy from the configuration.

Parameters *msap-policy-name* — Specifies the managed SAP policy name up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

create — Keyword used to create the managed SAP policy. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

cpu-protection

Syntax **cpu-protection**

Context config>sys>security
config>service>vprn>sub-if>grp-if>sap

Description This command enables the context to configure CPU protection policies.

cpu-protection

Syntax **cpu-protection** *policy-id* [**mac-monitoring**]
no cpu-protection

Context config>subscr-mgmt>msap-policy
config>service>ies>sub-if>grp-if>sap
config>service>vpls>sap
config>service>vprn>sub-if>grp-if>sap

Description This command assigns an existing CPU protection policy to the SAP or interface.

CPU protection policies are configured in the **config>sys>security>cpu-protection** context.

The **no** form of this command removes the policy ID from the SAP or interface configuration.

Default cpu-protection 254

Parameters *policy-id* — Specifies an existing CPU protection policy to assign to the SAP or interface.

Values 1 to 255

mac-monitoring — Specifies that the per-source rate limit be applied.

cpu-protection

Syntax **cpu-protection** *policy-id*
no cpu-protection

Context config>router>if
config>service>ies>if
config>service>vprn>if

Description This command assigns an existing CPU protection policy to the SAP or interface.

CPU protection policies are configured in the **config>sys>security>cpu-protection** context.

The **no** form of this command removes the policy ID from the SAP or interface configuration.

Default cpu-protection 254

Parameters *policy-id* — Specifies an existing CPU protection policy to assign to the SAP.

Values 1 to 255

default-host

Syntax **default-host** *ip-address*/*mask* **next-hop** *next-hop-ip*
no default-host

Context config>service>ies>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>sap

Description This command configures the default-host to be used. More than one default-host can be configured per SAP.

The **no** form of this command removes the values from the configuration.

Parameters *ip-address/mask* — Assigns an IP address/IP subnet format to the interface.

next-hop-ip — Assigns the next hop IP address.

dist-cpu-protection

Syntax **dist-cpu-protection** *policy-name*
no dist-cpu-protection

Context config>subscr-mgmt>msap-policy

Description This command assigns a Distributed CPU Protection (DCP) policy to an MSAP policy. The DCP policy is automatically assigned to MSAPs created with this policy. A non-existent DCP policy can be assigned to an **msap-policy** because an MSAP policy is similar to a template that is applied in the MSAP creation. The DCP policy is validated at the time that the MSAP is created, and the MSAP creation is blocked (and an appropriate log event created) if the DCP policy does not exist.

Note: For other types of objects (for example, normal non-MSAP SAPs and network interfaces) the DCP policy must exist before it can be assigned to the SAP.

The **no** form of this command removes the policy name from the configuration.

If no dist-cpu-protection policy is assigned to an MSAP policy, then the default access DCP policy (_default-access-policy) is used.

If no DCP functionality is required on the MSAP policy, then an empty DCP policy can be created and explicitly assigned to the MSAP policy.

Parameters *policy-name* — Specifies the name of the DCP policy up to 32 characters.

dynamic-services

Syntax **dynamic-services**

Context config>service>vpls>sap

Description This command enables the context to configure dynamic service policy parameters.

dynamic-services-policy

Syntax **dynamic-services-policy** *dynsvc-policy-name*
no dynamic-services-policy

Context config>service>vpls>sap

Description This command configures the dynamic service policy associated with this SAP.

ies-vprn-only-sap-parameters

Syntax **ies-vprn-only-sap-parameters**

Context config>subscr-mgmt>msap-policy

Description This command enables the context to configure managed SAP IES and VPRN properties. VPRN services are supported on the 7750 SR only.

egress

Syntax **egress**

Context config>subscr-mgmt>msap-policy>ies-vprn
config>subscr-mgmt>msap-policy>vpls-only

Description This command enables the context to configure egress policies for Managed SAPs (MSAPs).

ingress

Syntax **ingress**

Context config>subscr-mgmt>msap-policy>ies-vprn
config>subscr-mgmt>msap-policy>vpls-only

Description This command enables the context to configure ingress policies for Managed SAPs (MSAPs).

qos

Syntax **qos** *policy-id*
qos *policy-id* [**multipoint-shared** | **service-queuing**]
qos *policy-id* [**shared-queuing** | **service-queuing**]
no qos

Context config>subscr-mgmt>msap-policy>ies-vprn>egress
config>subscr-mgmt>msap-policy>ies-vprn>ingress
config>subscr-mgmt>msap-policy>vpls-only>egress
config>subscr-mgmt>msap-policy>vpls-only>ingress

Description This command specifies the ingress or egress Quality of Service (QoS) policy that is associated with a Managed SAP (MSAP). Only QoS policies with scope template can be associated with MSAPs.

The **no** form of this command resets the default value. The system default QoS policy is used in that case.

Default qos 1 — For egress parameters.

qos 1 multipoint-shared — For ingress vpls-only-sap-parameters.

qos 1 shared-queuing — For ingress ies-vprn-only-sap-parameters.

Parameters *policy-id* — Specifies the QoS policy ID or name to associate with the MSAPs. The policy ID or name must already exist.

Values id: 1 to 65535
name: policy name up to 64 characters.

multipoint-shared — Ingress unicast MSAP queues are mapped one-for-one with hardware queues. Multipoint MSAP queues are not instantiated. Unicast and BUM packets traverse the ingress forwarding plane twice: in the first pass, both unicast and BUM traffic use the unicast MSAP queues, while in the second pass shared unicast and multipoint queues are used towards the egress line card. Multipoint-shared queuing greatly reduces ingress queue consumption. This keyword can only be specified for ingress **vpls-only-sap-parameters**.

shared-queuing — Ingress unicast MSAP queues are mapped one-for-one with hardware queues. Unicast packets traverse the ingress forwarding plane twice: in the first pass the MSAP queues are used, while in the second pass shared queues are used towards the egress line card. Shared-queuing greatly reduces ingress queue consumption. This keyword can only be specified for ingress **ies-vprn-only-sap-parameters**.

service-queuing — Ingress MSAP queues are mapped to multiple hardware queues, per egress line card destination. For scaled environments such as in subscriber management, this could lead to a high number of queue resources being used. Therefore, service-queuing for MSAPs should only be enabled for specific use cases. This keyword can only be specified for ingress parameters.

igmp-host-tracking

Syntax **igmp-host-tracking**

Context config>subscr-mgmt>msap-policy

Description This command enables the context to configure IGMP host tracking parameters.

disable-router-alert-check

Syntax [**no**] **disable-router-alert-check**

Context config>service>vprn>sub-if>grp-if>sap>igmp-host-tracking

Description This command disables the IGMP router alert check option.

The **no** form of this command enables the router alert check.

expiry-time

Syntax **expiry-time** *expiry-time*
no expiry-time

Context config>subscr-mgmt>msap-policy>igmp-host-tracking

Description This command configures the time that the system continues to track inactive hosts.

The **no** form of this command removes the values from the configuration.

Parameters *expiry-time* — Specifies the time, in seconds, that this system continues to track an inactive host.

Values 1 to 65535

import

Syntax **import** *policy-name*
no import

Context config>subscr-mgmt>msap-policy>igmp-host-tracking

Description This command specifies the import routing policy to be used for IGMP packets to be used on this SAP or SDP. Only a single policy can be imported on a single SAP at any time.

The **no** form of this command removes the policy association from the SAP or SDP.

Parameters *policy-name* — Specifies the routing policy name up to 32 characters.

max-num-groups

Syntax **max-num-groups** *max-num-groups*
no max-num-groups

Context config>subscr-mgmt>msap-policy>igmp-host-tracking

Description This command configures the maximum number of multicast groups allowed to be tracked.

The **no** form of this command removes the values from the configuration.

Parameters *max-num-groups* — Specifies the maximum number of multicast groups allowed to be tracked.

Values 1 to 196607

max-num-sources

Syntax **max-num-sources** *max-num-sources*
no max-num-sources

Context config>subscr-mgmt>msap-policy>igmp-host-tracking

Description This command configures the maximum number of multicast sources allowed to be tracked per group.

The **no** form of this command removes the value from the configuration.

Parameters *max-num-sources* — Specifies the maximum number of multicast sources allowed to be tracked per group.

Values 1 to 1000

max-num-grp-sources

Syntax **max-num-grp-sources** [*max-num-grp-sources*]
no max-num-grp-sources

Context config>subscr-mgmt>msap-policy>igmp-host-tracking

Description This command configures the maximum number of group sources for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of group sources, the group sources that are already accepted are not deleted. Only new group sources will not be allowed. When this object has a value of 0, there is no limit to the number of group sources.

The **no** form of this command removes the value from the configuration.

Parameters *max-num-grp-sources* — Specifies the maximum number of multicast sources allowed to be tracked per group.

Values 1 to 32000

lag-link-map-profile

Syntax **lag-link-map-profile** *link-map-profile-id*
no lag-link-map-profile

Context config>subscr-mgmt>msap-policy

Description This command assigns a pre-configured lag link map profile to a SAP or network interface configured on a LAG or a PW port that exists on a LAG. Once assigned or de-assigned, the SAP or network interface egress traffic is re-hashed over LAG as required by the new configuration.

The **no** form of this command reverts the SAP/network interface to use per-flow, service or link hash as configured for the service/LAG.

Parameters *link-map-profile-id* — Defines a unique LAG link map profile on which the LAG the SAP/ network interface exist.

Default 1 to 64

sticky-msaps

Syntax [**no**] **sticky-msaps**

Context config>subscr-mgmt>msap-policy

Description This command prevents MSAPs associated with the specified MSAP policy from being deleted unless a manual **clear** command is issued. If this command is not enabled, an MSAP is deleted when a host creation fails or when a subscriber is no longer associated with the MSAP, for example, when a subscriber ends the session. This feature is useful for an operator who wants to keep historical statistics records on MSAPs. It can also speed up host creation on an MSAP since the MSAP is already created.

The **no** form of this command reverts to the default.

sub-sla-mgmt

Syntax [**no**] **sub-sla-mgmt**

Context config>service>vpls>sap
config>service>ies>sub-if>grp-if>sap
config>service>ies>if>sap
config>subscr-mgmt>msap-policy
config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt
config>service>vpls>sap>sub-sla-mgmt

Description This command enables the context to configure subscriber management parameters for a SAP or MSAP.

The **no** form of this command reverts to the default.

def-app-profile

Syntax **def-app-profile** *app-profile-name*
no def-app-profile

Context config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap-param>sub-sla-mgmt
config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt
config>service>vprn>sub-if>grp-if>sap-param>sub-sla-mgmt
config>service>vpls>sap>sub-sla-mgmt

Description This command specifies the application profile to be used by a subscriber host.

The **no** form of this command removes the application profile name from the configuration.

Parameters *app-profile-name* — Specifies an existing application profile to be mapped to the subscriber profile by default up to 32 characters.

def-inter-dest-id

Syntax **def-inter-dest-id** {**string** *string* | **use-top-q**}
no def-inter-dest-id

Context config>subscr-mgmt>msap-policy>sub-sla-mgmt

Description This command specifies a default destination string for all subscribers associated with the SAP. The command also accepts the **use-top-q** flag that automatically derives the string based on the top most delineating Dot1Q tag from the SAP's encapsulation.

The **no** form of this command removes the default subscriber identification string from the configuration.

Parameters **use-top-q** — Derives the string based on the top most delineating Dot1Q tag from the SAP's encapsulation.

string — Specifies the subscriber identification applicable for a subscriber host up to 32 characters.

def-sub-id

Syntax **def-sub-id use-auto-id**
def-sub-id use-sap-id
def-sub-id string *sub-id*
no def-sub-id

Context config>subscr-mgmt>msap-policy>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap-param>sub-sla-mgmt
config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt
config>service>vprn>sub-if>grp-if>sap-param>sub-sla-mgmt
config>service>vpls>sap>sub-sla-mgmt

Description This command specifies the explicit default sub-id for dynamic subscriber hosts (including ARP hosts) in case that the sub-id string is not supplied through RADIUS or LUDB.

The sub-id is assigned to a new subscriber host in the following order of priority:

- RADIUS
- LUDB
- Explicit default – The **def-sub-id** command is used to explicitly set the *sub-id* name of the host as one of the following:
  - The SAP ID to which the new host is associated
  - An explicit string
  - An auto-generated string consisting of the concatenated subscriber identification fields defined under the **subscr-mgmt**>**auto-sub-id-key** node. The fields are taken in the order in which they are configured and are separated by a '|'character. The subscriber host identification fields are separately defined for IPoE and PPPoE host types.
- Implicit default – If the sub-id string is not returned via RADIUS or LUDB and there is no def-sub-id configured, the *sub-id* name is generated as a random 10 character encoded string based on the auto-sub-id-keys. This 10 characters encoded string is unique per chassis as well as in dual-homed environment. It is generated based on auto-sub-id-keys. If auto-sub-id-keys are not explicitly configured, the default ones are:
  - mac, sap-id, or session-id for PPP type hosts
  - mac or sap-id for IPoE type hosts.

This command does not apply to static subscribers.

The **no** form of this command reverts to the default.

Parameters **use-sap-id** — Specifies the sub-id name-id on which the original request for host creation arrived (DHCP Discover, or PADI or ARP Request).

sub-id — Explicitly configured sub-id name up to 32 characters.

use-auto-id — Specifies the concatenated string of auto-sub-id-keys separated by a "|" character.

Default no def-sub-id

Implicit default If the *sub-id string* is not supplied through RADIUS, LUDB, or **def-sub-id** configuration, then a random 10 character encoded sub-id name is generated. This random sub-id name is based on the subscriber identification keys defined under the **subscr-mgmt**>**auto-sub-id-key** node. In case that the auto-sub-id-keys are not defined explicitly, the default ones are:

- mac, sap-id, or session-id for PPPoE type hosts
- mac or sap-id for IPoE type hosts

def-sla-profile

Syntax **def-sla-profile** *default-sla-profile-name*
no def-sla-profile

Context config>service>vpls>sap>sub-sla-mgmt
config>service>ies>if>sap>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt

Description This command specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sla-profile** context.

An SLA profile is a named group of QoS parameters used to define per service QoS for all subscriber hosts common to the same subscriber within a provider service offering. A single SLA profile may define the QoS parameters for multiple subscriber hosts. SLA profiles are maintained in two locations, the subscriber identification policy and the subscriber profile templates. After a subscriber host is associated with an SLA profile name, either the subscriber identification policy used to identify the subscriber or the subscriber profile associated with the subscriber host must contain an SLA profile with that name. If both the subscriber identification policy and the subscriber profile contain the SLA profile name, the SLA profile in the subscriber profile is used.

The **no** form of this command removes the default SLA profile from the SAP configuration.

Parameters *default-sla-profile-name* — Specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sla-profile** context.

def-sla-profile

Syntax **def-sla-profile** *default-sla-profile-name*
no def-sla-profile

Context config>subscr-mgmt>msap-policy>sub-sla-mgmt

Description This command specifies a default SLA profile for an MSAP.

An SLA profile is a named group of QoS parameters used to define per service QoS for all subscriber hosts common to the same subscriber within a provider service offering. A single SLA profile may define the QoS parameters for multiple subscriber hosts.

The **no** form of this command removes the default SLA profile from the MSAP configuration.

Parameters *default-sla-profile-name* — Specifies a default SLA profile for an MSAP up to 32 characters.

def-sub-profile

Syntax **def-sub-profile** *default-subscriber-profile-name*

Context config>subscr-mgmt>msap-policy>sub-sla-mgmt

Description This command specifies a default subscriber profile for an MSAP.

A subscriber profile defines the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for subscriber using the subscriber profile.

The **no** form of this command removes the default SLA profile from the SAP configuration.

Parameters *default-sub-profile* — Specifies a default subscriber profile for this SAP up to 32 characters.

multi-sub-sap

Syntax **multi-sub-sap** [**subscriber** *limit*]
no multi-sub-sap

Context config>subscr-mgmt>msap-policy>sub-sla-mgmt

Description This command defines the maximum number of subscribers (dynamic + static) that can be simultaneously active on an MSAP.

If the limit is reached, a new host is denied access and the corresponding DHCP ACK is dropped.

The **no** form of this command reverts back to the default setting.

Default multi-sub-sap 1

Parameters *limit* — Specifies the maximum number of subscribers allowed.

Note: The operational maximum value may be smaller due to equipped hardware dependencies.

Values 1 to 131071

single-sub-parameters

Syntax **single-sub-parameters**

Context config>subscr-mgmt>msap-policy>sub-sla-mgmt

Description This command enables the context to configure single subscriber MSAP parameters.

non-sub-traffic

Syntax **non-sub-traffic sub-profile** *sub-profile-name* **sla-profile** *sla-profile-name* [**subscriber** *sub-ident-string*] [**app-profile** *app-profile-name*]
no non-sub-traffic

Context config>subscr-mgmt>msap-policy>sub-sla-mgmt>single-sub

Description This command configures traffic profiles for non-IP traffic such as PPPoE. It is used in conjunction with the profiled-traffic-only on single subscriber SAPs and creates a subscriber host which is used to forward non-IP traffic through the single subscriber SAP without the need for SAP queues.

The **no** form of this command removes any configured profile.

Parameters *sub-profile-name* — Identifies the subscriber profile name up to 32 characters.

sla-profile-name — Identifies the SLA profile name up to 32 characters.

profiled-traffic-only

Syntax [**no**] **profiled-traffic-only**

Context config>subscr-mgmt>msap-policy>sub-sla-mgmt>single-sub

Description This command specifies whether only profiled traffic is applicable for an MSAP. When enabled, all queues is deleted.

The **no** form of this command reverts to the default setting.

sub-ident-policy

Syntax [**no**] **sub-ident-policy** *sub-ident-policy-name*

Context config>subscr-mgmt>msap-policy>sub-sla-mgmt

Description This command specifies an existing subscriber identification policy. Each subscriber identification policy can have a default subscriber profile defined. The subscriber identification policy default subscriber profile overrides the system default and the subscriber SAP default subscriber profiles. Defining a subscriber identification policy default subscriber profile is optional.

Defining a subscriber profile as a subscriber identification policy default subscriber profile will cause all active subscribers currently associated with a subscriber SAP using the policy and associated with a subscriber policy through the system default or subscriber SAP default subscriber profiles to be reassigned to the subscriber policy defined as default on the subscriber identification policy.

Attempting to delete a subscriber profile that is currently defined as a default for a subscriber identification policy will fail.

When attempting to remove a subscriber identification policy default subscriber profile definition, the system will evaluate each active subscriber on all subscriber SAPs the subscriber identification policy is currently associated with that are using the default definition to determine whether the active subscriber can be either reassigned to a subscriber SAP default or the system default subscriber profile. If all active subscribers cannot be reassigned, the removal attempt will fail.

The **no** form of this command reverts to the default.

Parameters *sub-ident-policy-name* — Specifies the name of the subscriber identification policy up to 32 characters.

vpls-only-sap-parameters

Syntax **vpls-only-sap-parameters**

Context config>subscr-mgmt>msap-policy

Description This command enables the context to configure MSAP VPLS properties.

min-auth-interval

Syntax **min-auth-interval** *min-auth-interval*
no min-auth-interval

Context config>subscr-mgmt>msap-policy>vpls-only
config>service>vpls>sap>arp-host
config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if>arp-host

Description This command configures the minimum authentication interval.

The **no** form of this command reverts to the default.

Default min-auth-interval 15

Parameters *min-auth-interval* — Specifies the minimum authentication interval.

Values 1 to 6000

sap-host-limit

Syntax **sap-host-limit** *max-num-hosts-sap*
no sap-host-limit

Context config>service>ies>sub-if>grp-if>arp-host
config>service>vprn>sub-if>grp-if>arp-host

Description This command configures the maximum number of ARP hosts per SAP.

The **no** form of this command reverts to the default.

Default sap-host-limit 1

Parameters *max-num-hosts-sap* — Specifies the maximum number of ARP hosts per SAP allowed on this interface.

Note: The operational maximum value may be smaller due to equipped hardware dependencies.

Values 1 to 131071

arp-reply-agent

Syntax **arp-reply-agent** [**sub-ident**]
no arp-reply-agent

Context config>subscr-mgmt>msap-policy>vpls-only

Description This command enables a special ARP response mechanism in the system for ARP requests destined to static or dynamic hosts associated with the SAP. The system responds to each ARP request using the hosts MAC address as the both the source MAC address in the Ethernet header and the target hardware address in the ARP header.

ARP replies and requests received on an MSAP with **arp-reply-agent** enabled is evaluated by the system against the anti-spoof filter entries associated with the ingress SAP (if the SAP has anti-spoof filtering enabled). ARPs from unknown hosts on the SAP is discarded when anti-spoof filtering is enabled.

The ARP reply agent only responds if the ARP request enters an interface (SAP, spoke-SDP or mesh-SDP) associated with the VPLS instance of the MSAP.

A received ARP request that is not in the ARP reply agent table is flooded to all forwarding interfaces of the VPLS capable of broadcast except the ingress interface while honoring split-horizon constraints.

Static hosts can be defined using the **host** command. Dynamic hosts are enabled on the system by enabling the **lease-populate** command in the **dhcp** context. In the event that both a static host and a dynamic host share the same IP and MAC address, the VPLS ARP reply agent will retain the host information until both the static and dynamic information are removed. In the event that both a static and dynamic host share the same IP address, but different MAC addresses, the VPLS ARP reply agent is populated with the static host information.

The **arp-reply-agent** command will fail if an existing static host does not have both MAC and IP addresses specified. Once the ARP reply agent is enabled, creating a static host on the MSAP without both an IP address and MAC address will fail.

The ARP-reply-agent may only be enabled on SAPs supporting Ethernet encapsulation.

The **no** form of this command disables ARP-reply-agent functions for static and dynamic hosts on the MSAP.

Parameters **sub-ident** — Configures the arp-reply-agent to discard ARP requests received on the MSAP that are targeted for a known host on the same MSAP with the same subscriber identification.

Hosts are identified by their subscriber information. For DHCP subscriber hosts, the subscriber hosts, the subscriber information is configured using the optional subscriber parameter string.

When arp-reply-agent is enabled with **sub-ident**:

- If the subscriber information for the destination host exactly matches the subscriber information for the originating host and the destination host is known on the same MSAP as the source, the ARP request is silently discarded.
- If the subscriber information for the destination host or originating host is unknown or undefined, the source and destination hosts are not considered to be the same subscriber. The ARP request is forwarded outside the MSAP's Split Horizon Group.
- When **sub-ident** is not configured, the arp-reply-agent does not attempt to identify the subscriber information for the destination or originating host and will not discard an ARP request based on subscriber information.

dhcp

Syntax **dhcp**

Context config>subscr-mgmt>msap-policy>vpls-only

Description This command enables the context to configure DHCP parameters.

circuit-id

Syntax **circuit-id** [**ascii-tuple** | **vlan-ascii-tuple**]

no circuit-id

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option

Description When enabled, the router sends an ASCII-encoded tuple in the **circuit-id** sub-option of the DHCP packet. This ASCII-tuple consists of the access-node-identifier, service-id, and SAP-ID, separated by "|".

If disabled, the **circuit-id** sub-option of the DHCP packet is left empty.

The **no** form of this command reverts to the default.

Default circuit-id ascii-tuple

Parameters **ascii-tuple** — Specifies that the ASCII-encoded concatenated tuple consisting of the access-node-identifier, service-id, and interface-name is used.

vlan-ascii-tuple — Specifies that the format will include VLAN-id and dot1p bits in addition to what is included in ascii-tuple already. The format is supported on dot1q and qinq ports only. Thus, when the option 82 bits are stripped, dot1p bits is copied to the Ethernet header of an outgoing packet.

remote-id

Syntax **remote-id** [**mac** | **string** *string*]
no remote-id

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option

Description This command specifies what information goes into the remote-id sub-option in the DHCP relay packet.

If disabled, the **remote-id** sub-option of the DHCP packet will be left empty.

The **no** form of this command reverts to the default.

Parameters **mac** — Specifies the MAC address of the remote end is encoded in the sub-option.

string — Specifies the remote-id up to 32 characters.

vendor-specific-option

Syntax [**no**] **vendor-specific-option**

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option

Description This command configures the vendor-specific sub-option parameters of the DHCP relay packet.

The **no** form of this command reverts to the default.

client-mac-address

Syntax [**no**] **client-mac-address**

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor

Description This command enables the sending of the MAC address in the Nokia vendor-specific sub-option of the DHCP relay packet.

The **no** form of this command disables the sending of the MAC address in the Nokia vendor-specific sub-option of the DHCP relay packet.

sap-id

Syntax [**no**] **sap-id**

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor
config>service>ies>sub-if>dhcp>option

Description This command enables the sending of the SAP ID in the A Nokia vendor-specific sub-option of the DHCP relay packet.

The **no** form of this command disables the sending of the SAP ID in the Nokia vendor-specific sub-option of the DHCP relay packet.

service-id

Syntax [**no**] **service-id**

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor
config>service>ies>sub-if>dhcp>option

Description This command enables the sending of the service ID in the Nokia vendor-specific sub-option of the DHCP relay packet.

The **no** form of this command disables the sending of the service ID in the Nokia vendor-specific sub-option of the DHCP relay packet.

string

Syntax [**no**] **string** *text*

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor
config>service>ies>sub-if>dhcp>option

Description This command specifies the string in the Nokia vendor-specific sub-option of the DHCP relay packet.

The **no** form of this command returns the default value.

Parameters *text* — The string can be any combination of ASCII characters up to 32 characters. If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

system-id

Syntax [**no**] **system-id**

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>option>vendor

Description This command specifies whether the system-id is encoded in the A Nokia vendor-specific sub-option of Option 82.

The **no** form of this command reverts to the default.

lease-time

Syntax **lease-time** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*] [**radius-override**]
no lease-time

Context config>subscr-mgmt>msap-policy>vpls-only>dhcp>proxy

Description This command defines the length of lease-time that is provided to DHCP clients. By default the local-proxy-server will always make use of the lease-time information provide by either a RADIUS or DHCP server.

The **no** form of this command disables the use of the lease-time command. The local-proxy-server will use the lease-time offered by either a RADIUS or DHCP server.

Parameters **radius-override** — Specifies that the local-proxy-server will use the configured lease-time information to provide DHCP clients. This is supported only in the **config>service>vpls>sap>dhcp>proxy-server** context.

days — Specifies the number of days that the given IP address is valid.

Values 0 to 3650

hours — Specifies the number of hours that the given IP address is valid.

Values 0 to 23

minutes — Specifies the number of minutes that the given IP address is valid.

Values 0 to 59

seconds — Specifies the number of seconds that the given IP address is valid.

Values 0 to 59

igmp-snooping

Syntax **igmp-snooping**

Context config>subscr-mgmt>msap-policy>vpls-only

Description This command enables the context to configure Internet Group Management Protocol (IGMP) snooping parameters.

fast-leave

Syntax [**no**] **fast-leave**

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

Description This command enables fast leave.

When IGMP fast leave processing is enabled, the 7450 ESS or 7750 SR will immediately remove a SAP or SDP from the IP multicast group when it detects an IGMP leave on that SAP or SDP. Fast leave processing allows the switch to remove a SAP or SDP that sends a leave from the forwarding table without first sending out group-specific queries to the SAP or SDP, and thus speeds up the process of changing channels (zapping).

Fast leave should only be enabled when there is a single receiver present on the SAP or SDP.

When fast leave is enabled, the configured **last-member-query-interval** value is ignored.

The **no** form of this command reverts to the default.

import

Syntax **import** *policy-name*
no import

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

Description This command specifies the import routing policy to be used for IGMP packets to be used on this SAP or SDP. Only a single policy can be imported on a single SAP at any time.

The **no** form of this command removes the policy association from the SAP or SDP.

Parameters *policy-name* — Specifies the routing policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes. Routing policies are configured in the **config>router>policy-options** context The router policy must be defined before it can be imported.

last-member-query-interval

Syntax **last-member-query-interval** *tenths-of-seconds*
no last-member-query-interval

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

Description This command configures the maximum response time used in group-specific queries sent in response to leave messages, and is also the amount of time between two consecutive group-specific queries. This value may be tuned to modify the leave latency of the network. A reduced value results in reduced time to detect the loss of the last member of a group.

The configured interval is ignored when fast-leave is enabled on the SAP or SDP.

The **no** form of this command reverts to the default.

Default last-member-query-interval 10

Parameters *seconds* — Specifies the frequency, in tenths of seconds, at which query messages are sent.

Values 1 to 50

max-num-groups

Syntax **max-num-groups** *max-num-groups*
no max-num-groups

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

Description This command configures the maximum number of multicast groups that can be joined on an MSAP or SDP. If the router receives an IGMP join message that would exceed the configured number of groups, the request is ignored.

Parameters *max-num-groups* — Specifies the maximum number of groups that can be joined on an MSAP or SDP.

Values 1 to 1000

mcac

Syntax **mcac**

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>

Description This command enables the context to configure multicast CAC parameters.

if-policy

Syntax **if-policy** *mcac-if-policy-name*
no if-policy

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac

Description This command assigns an existing MCAC interface policy to this MSAP policy.

The **no** form of this command removes the MCAC interface policy association.

Parameters *mcac-if-policy-name* — Specifies an existing MCAC interface policy up to 32 characters.

mc-constraints

Syntax **mc-constraints**

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac

Description This command enables the context to configure the level and its associated bandwidth for a bundle or a logical interface.

level

Syntax **level** *level-id* **bw** *bandwidth*
no level *level-id*

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac

Description This command configures levels and their associated bandwidth for multicast CAC policy on an interface.

The **no** form of this command reverts to the default.

Parameters *level-id* — Specifies has an entry for each multicast CAC policy constraint level configured on a system.

Values 1 to 8

bandwidth — Specifies the bandwidth in kilobits per second (kb/s) for the level.

Values 1 to 2147483647

number-down

Syntax **number-down** *number-lag-port-down* **level** *level-id*
no number-down *number-lag-port-down*

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac

Description This command configures the number of ports down along with level for multicast CAC policy on an MSAP

The **no** form of this command reverts to the default.

Parameters *number-lag-port-down* — Specifies the number of port in a LAG group that are down. If the number of ports available in the LAG is reduced by the number of ports configured in this command here then bandwidth allowed for bundle and/or interface is as per the levels configured in this context.

Values 1 to 64 (for 64-link LAG)
1 to 32 (for other LAGs)

level-id — Specifies the amount of bandwidth available within a given bundle for MC traffic for a specified level.

policy

Syntax **policy** *policy-name*
no policy

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac

Description This command configures the multicast CAC policy name.

The **no** form of this command reverts to the default.

Parameters *policy-name* — The multicast CAC policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

unconstrained-bw

Syntax **unconstrained-bw** *bandwidth* **mandatory-bw** *mandatory-bw*
no unconstrained-bw

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac

Description This command configures the bandwidth for the interface's multicast CAC policy traffic. When disabled (**no unconstrained-bw**) there is no checking of bandwidth constraints on the interface level. When enabled and a policy is defined, enforcement is performed. The allocated bandwidth for optional channels should not exceed the **unconstrained-bw** minus the **mandatory-bw** and the mandatory channels have to stay below the specified value for the **mandatory-bw**. After this interface check, the bundle checks are performed.

The **no** form of this command reverts to the default.

Parameters *bandwidth* — Specifies bandwidth assigned for interface's MCAC policy traffic, in kilobits per second (kb/s).

Values 0 to 2147483647

mandatory-bw — Specifies the bandwidth pre-reserved for all the mandatory channels on a given interface in kilobits per second (kb/s).

If the *bandwidth* value is 0, no mandatory channels are allowed. If the value of *bandwidth* is '-1', then all mandatory and optional channels are allowed.

If the value of *mandatory-bw* is equal to the value of *bandwidth*, then all the unconstrained bandwidth on a given interface is allocated to mandatory channels configured through multicast CAC policy on that interface and no optional groups (channels) are allowed.

The value of *mandatory-bw* should always be less than or equal to that of *bandwidth*, An attempt to set the value of *mandatory-bw* greater than that of *bandwidth*, will result in inconsistent value error.

Values 0 to 2147483647

use-lag-port-weight

Syntax [**no**] **use-lag-port-weight**

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mcac>mc-constraints

Description This command enables port weight to be used when determining available bandwidth per level when LAG ports go down/come up. The command is required for proper operation on mixed port-speed LAGs and can be used for non-mixed port-speed LAGs as well.

The **no** form of this command reverts to the default.

sub-mcac-policy

Syntax **sub-mcac-policy** *sub-mcac-policy-name* [**create**]
no sub-mcac-policy *b*

Context config>subscr-mgmt

Description This command creates a policy template with MCAC bandwidth limits that are applied to the subscriber.

Per interface mcac bandwidth limits are set directly under the interface (regular interface or group-interface) and no such policy templates are needed.

The need for a separate policy template for subscribers is due to the fact that groups of subscribers under the same group-interface can share certain settings that can be configured via this template.

To summarize, the MCAC bandwidth constraints for subscribers are defined in the sub-mcac-policy while the mcac bandwidth constraints for the interface are configured directly under the **igmp>interface>mcac** or **igmp>grp-if>mcac** context without the need for policy templates.

Note: The sub-mcac-policy only deals with the mcac bandwidth limits and not the channel bandwidth definitions. Channels bandwidth is defined in a different policy (in the **config>router>mcac** context) and that policy is applied on the interface level as follows:

- For group-interface: under the **config>service>vprn>igmp>grp-if>mcac** context
- For regular interface: under the **config>service/router>igmp>interface>mcac** context.

In case of HQoS Adjustment, it is mandatory that the sub-mcac-policy be created and applied to the subscriber. The sub-mac-policy does not have to contain any bandwidth constrains, but it has to be in a no shutdown state in order for HQoS Adjustment to work.

The **no** form of this command reverts to the default.

Parameters *policy-name* — Specifies the name of the policy up to 32 characters.

mvr

Syntax **mvr**

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

Description This command enables the context to configure Multicast VPLS Registration (MVR) parameters.

from-vpls

Syntax **from-vpls** *service-id*
no from-vpls

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp>mvr

Description This command configures the VPLS from which multicast traffic is copied upon receipt of an IGMP join request.

IGMP snooping must be enabled on the MVR VPLS.

The **no** form of this command reverts to the default.

Parameters *service-id* — Specifies the MVR VPLS from which multicast channels be copied into an MSAP.

Values
service-id: 1 to 2147483647
service-name: up to 64 characters (applies only to the 7750 SR)

query-interval

Syntax **query-interval** *seconds*
no query-interval

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

Description This command configures the IGMP query interval. If the **send-queries** command is enabled, this parameter specifies the interval between two consecutive general queries sent by the system on an MSAP or SDP.

The configured query interval must be greater than the configured **query-response** interval.

If **send-queries** is not enabled on an MSAP or SDP, the configured query interval value is ignored.

Default query-interval 125

Parameters *seconds* — Specifies the time interval, in seconds, that the router transmits general host-query messages

Values 2 to 1024

query-response-interval

Syntax **query-response-interval** *seconds*

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

Description This command configures the IGMP query response interval. If the **send-queries** command is enabled, this parameter specifies the maximum response time advertised in IGMPv2/v3 queries.

The configured query-response-interval must be smaller than the configured **query-interval**.

If **send-queries** is not enabled on an MSAP or SDP, the configured query-response-interval value is ignored.

The **no** form of this command reverts to the default.

Default query-response-interval 10

Parameters *seconds* — Specifies the length of time, in seconds, to wait to receive a response to the host-query message from the host.

Values 1 to 1023

robust-count

Syntax **robust-count** *robust-count*
no robust-count

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

Description This command configures the IGMP robustness variable. If the **send-queries** command is enabled, this parameter allows tuning for the expected packet loss on a SAP or SDP. The robust-count variable allows tuning for the expected packet loss on a subnet and is comparable to a retry count. If an MSAP or SDP is expected to be “lossy”, this parameter may be increased. IGMP snooping on an MSAP or SDP is robust to (robust-count-1) packet losses.

If send-queries is not enabled, this parameter is ignored.

The **no** form of this command reverts to the default.

Default robust-count 2

Parameters *robust-count* — Specifies the robust count for the SAP or SDP.

Values 2 to 7

send-queries

Syntax [**no**] **send-queries**

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

Description This command specifies whether to send IGMP general query messages on the managed SAP. When send-queries is configured, all type of queries generate ourselves are of the configured version. If a report of a version higher than the configured version is received, the report will get dropped and a new wrong version counter will get incremented.

If send-queries is not configured, the version command has no effect. The version used on that SAP/SDP is the version of the querier. This implies that, for example, when there is a v2 querier, a v3 group or group-source specific query is never sent when a host wants to leave a certain group.

The **no** form of this command reverts to the default.

version

Syntax **version** *version*
no version

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

Description This command specifies the version of IGMP which is running on an MSAP. This object can be used to configure a router capable of running either value. For IGMP to function correctly, all routers on a LAN must be configured to run the same version of IGMP on that LAN.

When the **send-query** command is configured, all type of queries generate ourselves are of the configured **version**. If a report of a version higher than the configured version is received, the report gets dropped and a new “wrong version” counter is incremented.

If the **send-query** command is not configured, the **version** command has no effect. The version used on that SAP or SDP is the version of the querier. This implies that, for example, when there is a v2 querier, a v3 group or group-source specific query when a host wants to leave a certain group will never be sent.

Default version 3

Parameters *version* — Specifies the IGMP version.

Values 1, 2, 3

mac-da-hashing

Syntax [**no**] **mac-da-hashing**

Context config>subscr-mgmt>msap-policy>vpls-only

Description This command specifies whether subscriber traffic egressing a LAG SAP has its egress LAG link selected by a function of the MAC destination address instead of the subscriber ID.

This command is only meaningful if subscriber management is enabled and can be configured for a VPLS service.

The **no** form of this command reverts to the default.

split-horizon-group

Syntax **split-horizon-group** *group-name*

Context config>subscr-mgmt>msap-policy>vpls-only>igmp-snp

Description This command specifies the name of the split horizon group to which the MSAP belongs.

Parameters *group-name* — Specifies the split horizon group name up to 32 characters.

default-msap-policy

Syntax **default-msap-policy** *policy-name*
no default-msap-policy

Context config>service>vpls>sap

Description This command specifies the default managed SAP policy to use to create MSAPs when the response from the RADIUS server does not specify a managed SAP policy.

The *policy-name* parameter is only valid for a SAP with the keywords **capture-sap** specified in the SAP's configuration. The **capture-sap** keyword in the SAP configuration captures the SAP where triggering packets is sent to the CPM. Non-triggering packets captured by the capture SAP is dropped.

The managed SAP policy must already be defined in the **config>subscr-mgmt>msap-policy** context

The **no** form of this command removes the policy-name from the configuration.

Parameters *policy-name* — Specifies an existing default managed SAP policy.

trigger-packet

Syntax **trigger-packet** [**dhcp**] [**pppoe**] [**arp**] [**dhcp6**] [**ppp**] [**rtr-solicit**] [**data**]
no trigger-packet

Context config>service>vpls>sap

Description This command enables triggering packet to initiate RADIUS authentication that provides a service context. The authentication, together with the service context for this request, creates a managed SAP. The VLAN is the same as the triggering packet. This SAP behaves as a regular SAP but the configuration is not user-editable and not maintained in the configuration file. The managed SAP remains active as long as the session is active.

The **no** form of this command reverts to the default.

Parameters **dhcp** — Specifies whether the receipt of DHCP trigger packets on this VPLS SAP when the keyword **capture-sap** is specified in the **sap** command creation string, will result in a RADIUS authentication that will provide a service context and the creation of a SAP with a value of managed.

pppoe — Specifies whether the receipt of PPPoE trigger packets on this VPLS SAP when the keyword **capture-sap** is specified in the **sap** command creation string, will result in a RADIUS authentication that will provide a service context and the creation of a SAP with a value of managed.

arp — Indicates that ARP is the type of trigger packets for this entry.

dhcp6 — Indicates that DHCP6 is the type of trigger packets for this entry.

ppp — Indicates that PPP is the type of trigger packets for this entry.

rtr-solicit — Indicates that router solicit is the type of trigger packets for this entry.

data — Indicates that data is the type of trigger packets for this entry.

allow-dot1q-msaps

Syntax [**no**] **allow-dot1q-msaps**

Context config>service>vpls>sap

Description This command enables support for single tagged traffic triggering managed SAP creation on a qinq encapsulated capture SAP.

With this command enabled, a single tagged trigger packet received on a qinq encapsulated capture SAP (x/y/z:*.* or x/y/z:tag.*) can trigger the creation of an x/y/z:tag.0 managed SAP (MSAP).

The **config>system>ethernet>new-qinq-untagged-sap** command should be configured:

- as a prerequisite for an x/y/z:tag.* capture-sap
- where x/y/z:tag1.0 and x/y/z:tag1.tag2 MSAPs for an x/y/z:*.* capture-sap should co-exist

Note that enabling new-qinq-untagged-sap affects the behavior of existing <port-id>:tag.0 SAPs.

With the allow-dot1q-msaps command disabled (default), a single tagged trigger packet received on a qinq encapsulated capture SAP (x/y/z:*.* or x/y/z:tag.*) is dropped as “Invalid QTag”.

This command cannot be enabled on:

- a dot1q encapsulated **capture-sap**
- an inverse capture sap (x/y/z:*.tag)

The **no** form of this command reverts to the default.

authentication-policy

Syntax **authentication-policy** *name*
no authentication-policy

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command defines which subscriber authentication policy must be applied when a DHCP message is received on the interface. The authentication policies must already be defined. The policy will only be applied when DHCP snooping is enabled on the SAP on Layer 2 interfaces.

The **no** form of this command reverts to the default.

Parameters *name* — Specifies a unique authentication policy name.

msap-defaults

Syntax **msap-default**

Context config>subscr-mgmt>loc-user-db>ipoe>host
config>subscr-mgmt>loc-user-db>ppp>host

Description This command configures MSAP authentication defaults.

group-interface

Syntax **group-interface** *ip-int-name* [**prefix** {*port-id*}] [**suffix** {*port-id*}]
no group-interface

Context config>subscr-mgmt>loc-user-db>ipoe>host>msap-defaults
config>subscr-mgmt>loc-user-db>ppp>host>msap-defaults

Description This command configures the group interface.

The **no** form of this command removes the group interface parameters from the configuration.

Parameters *ip-int-name* — Specifies the IP interface name up to 32 characters.

prefix *port-id* — Specifies the port ID as the prefix to the specified IP interface name.

suffix *port-id* — Specifies the port ID as the suffix to the specified IP interface name.

policy

Syntax **policy** *msap-policy-name*
no policy

Context config>subscr-mgmt
config>subscr-mgmt>loc-user-db>ipoe>host>msap-defaults
config>subscr-mgmt>loc-user-db>ppp>host>msap-defaults

Description This command configures the MSAP policy.

The **no** form of this command removes the MSAP policy name from the configuration.

Parameters *msap-policy-name* — Specifies the policy name up to 32 characters.

service

Syntax **service** *service-id*
no service

Context config>subscr-mgmt>loc-user-db>ipoe>host>msap-defaults
config>subscr-mgmt>loc-user-db>ppp>host>msap-defaults

Description This command sets default service for all subscribers created based on trigger packets received on the given capture SAP in case the corresponding VSA is not included in the RADIUS authentication response. This command is applicable to capture SAP only.

The **no** form of this command reverts to the default.

Parameters *service-id* — Specifies the service ID as an integer or a name.

Values *service-id* - 1 to 2147483648
service-name - up to 64 characters

9.25.2.1.14 Multicast Listener Discovery (MLD) Commands

mld-policy

Syntax **mld-policy** *mld-policy-name* [**create**]
no mld-policy *mld-policy-name*

Context config>subscr-mgmt

Description This command enables the context to create an MLD policy.

The **no** form of this command reverts to the default.

Parameters *mld-policy-name* — Specifies the MLD policy name up to 32 characters.

create — Keyword required to create the configuration context.

disable-router-alert-check

Syntax [**no**] **disable-router-alert-check**

Context config>subscr-mgmt>mld-policy

Description This command disables router alert checking for MLD messages received on this interface.

The **no** form of this command enables router alert checking.

egress-rate-modify

Syntax **egress-rate-modify agg-rate-limit**
egress-rate-modify scheduler *scheduler-name*
no egress-rate-modify

Context config>subscr-mgmt>mld-policy

Description This command configures the egress rate modification.

The **no** form of this command removes the values from the configuration.

Parameters **agg-rate-limit** — Specifies that the maximum total rate for all subscriber egress queues for each subscriber associated with the policy.

scheduler-name — Specifies the scheduler to be applied for egress rate modification.

fast-leave

Syntax [**no**] **fast-leave**

Context config>subscr-mgmt>mld-policy

Description This command enables fast leave. When fast leave processing is enabled, the router immediately removes a SAP or SDP from the IP multicast group when it detects an MLD leave on that SAP or SDP. Fast leave processing allows the switch to remove a SAP or SDP that sends a leave from the forwarding table without first sending out group-specific queries to the SAP or SDP, and thus speeds up the process of changing channels (zapping).

Fast leave should only be enabled when there is a single receiver present on the SAP or SDP.

When fast leave is enabled, the configured last-member-query-interval value is ignored.

The **no** form of this command reverts to the default.

Default fast-leave

import

Syntax **import** *policy-name*
no import

Context config>subscr-mgmt>mld-policy

Description This command specifies the import routing policy to be used. Only a single policy can be imported at a time.

The **no** form of this command removes the policy association.

Parameters *policy-name* — Specifies the import policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes. Routing policies are configured in the **config>router>policy-options** context The router policy must be defined before it can be imported.

max-num-groups

Syntax **max-num-groups** *count*
no max-num-groups

Context config>subscr-mgmt>mld-policy

Description This command defines the maximum number of multicast groups that can be joined. If the router receives a join message that would exceed the configured number of groups, the request is ignored.

The **no** form of this command reverts to the default.

Parameters *count* — Specifies the maximum number of groups that can be joined.

Values 1 to 1000

max-num-grp-sources

Syntax **max-num-grp-sources** [*max-num-grp-sources*]
no max-num-grp-sources

Context config>subscr-mgmt>mld-policy

Description This command configures the maximum number of group sources for which MLD can have local receiver information based on received MLD reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of group sources, the group sources that are already accepted are not deleted. Only new group sources are not allowed. When this object has a value of 0, there is no limit to the number of group sources.

The **no** form of this command removes the value from the configuration.

Parameters *max-num-grp-sources* — Specifies the maximum number of multicast sources allowed to be tracked per group

Values 1 to 32000

max-num-sources

Syntax **max-num-sources** *max-num-sources*
no max-num-sources

Context config>subscr-mgmt>mld-policy

Description This command configures the maximum number of multicast sources allowed per group.

The **no** form of this command removes the value from the configuration.

Parameters *max-num-sources* — Specifies the maximum number of multicast sources allowed per group.

Values 1 to 1000

per-host-replication

Syntax [**no**] **per-host-replication** [**uni-mac** | **mcast-mac**]

Context config>subscr-mgmt>mld-policy

Description This command enables per-host-replication. In the per-host-replication mode, multicast traffic is replicated per each host within the subscriber irrespective of the fact that some hosts may be subscribed to the same multicast stream. As a result, in case that multiple hosts within the subscriber are registered for the same multicast group, the multicast streams of that group are generated. The destination MAC address of multicast streams is changed to unicast so that each host receives its own copy of the stream. Multicast traffic in the per-host-replication mode can be classified via the existing QoS CLI structure. As such the multicast traffic flows through the subscriber queues. HQoS Adjustment is not needed in this case.

The alternative behavior for multicast replication in IPoE environment is per-SAP- replication. In this model, only a single copy of the multicast stream is sent per SAP, irrespective of the number of hosts that are subscribed to the same multicast group. This behavior applies to 1:1 connectivity model as well as on 1:N connectivity model (SAP centric behavior as opposed to subscriber centric behavior).

In the per-SAP-replication model the destination MAC address is multicast (as opposed to unicast in the per-host-replication model). Multicast traffic is flowing via the SAP queue which is outside of the subscriber context. The consequence is that multicast traffic is not accounted in the subscriber HQoS. In addition, HQoS Adaptation is not supported in the per SAP replication model.

The **no** form of this command reverts to the default.

Parameters **uni-mac** — Specifies that multicast traffic s sent with a unicast MAC and multicast IP.

mcast-mac — Specifies the multicast traffic is sent with a multicast MAC and IP.

query-interval

Syntax **query-interval** *seconds*
no query-interval

Context config>subscr-mgmt>mld-policy

Description This command specifies the frequency at which the querier router transmits general host-query messages. Host-query messages solicit group membership information and are sent to the link-scope all-node address, FF02::1. If nothing is configured, by default, the **query-interval** takes the value defined in the **config**>**router**>**mld** context or in the **config**>**service**>**vprn**>**mld** context. It is highly recommended that all three query intervals be configured together on each MLD policy.

The **no** form of this command reverts to the default.

Parameters *seconds* — Specifies the frequency, in seconds, at which the router transmits general host-query messages.

Values 2 to 1024

query-last-listener-interval

Syntax **query-last-listener-interval** *seconds*
no query-last-listener-interval

Context config>subscr-mgmt>mld-policy

Description This command configures the frequency at which the querier router sends a group-specific query messages, including the messages sent in response to leave-group messages. The shorter the interval, the faster the loss of the last member of a group can be detected. If nothing is configured, by default, the **query-last-listener-interval** takes the value defined in the **config**>**router**>**mld** context or in the **config**>**service**>**vprn**>**mld** context. It is highly recommended that all three query intervals be configured together on each MLD policy.

The **no** form of this command reverts to the default.

Parameters *seconds* — Specifies the frequency, in seconds, at which query messages are sent.

Values 1 to 1023

query-response-interval

Syntax **query-response-interval** *seconds*
no query-response-interval

Context config>subscr-mgmt>mld-policy

Description This command configures the query response interval. If nothing is configured, by default, the **query-response-interval** takes the value defined in the **config**>**router**>**mld** context or in the **config**>**service**>**vprn**>**mld** context. It is highly recommended that all three query intervals be configured together on each MLD policy.

The **no** form of this command reverts to the default.

Parameters *seconds* — Specifies the length of time that the querier router waits to receive a response to from the host.

Values 1 to 1023

redirection-policy

Syntax **redirection-policy** *policy-name*
no redirection-policy

Context config>subscr-mgmt>mld-policy

Description This command applies multicast redirection action to the subscriber. The redirection action along with the redirected interface (and possibly service id) is defined in the referenced policy-name. MLD messages is redirected to an alternate interface if that alternate interface has MLD enabled. The alternate interface does not have to have any multicast groups registered via MLD. Currently all MLD messages are redirected and there is no ability to selectively redirect MLD messages based on match conditions (multicast-group address, source IP address, and so on). Multicast redirection is supported between VPRN services and also between interfaces within the Global Routing Context. Multicast Redirection is not supported between the VPRN services and the Global Routing Table (GRT).

MLD state is maintained per subscriber host and per redirected interface. Traffic is however forwarded only on the redirected interface.

The **no** form of this command removes the policy name from the configuration.

Parameters *policy-name* — Specifies a redirection policy name up to 32 characters. This is a regular policy defined under the **configure>router>policy-option>policy-statement** context.

static

Syntax **static**

Context config>subscr-mgmt>mld-policy

Description This command enables the context to configure MLD static group membership parameters.

group

Syntax [**no**] **group** *grp-ipv6-address*

Context config>subscr-mgmt>mld-policy>static

Description This command configures a static multicast group.

The **no** form of this command reverts to the default.

Parameters *grp-ipv6-address* — Specifies the IPv6 address.

Values <grp-ipv6-address>: ipv6-address - x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d
x - [0 to FFFF]H
d - [0 to 255]D
- multicast group IPv6 address

source

Syntax [**no**] **source** *ipv6-address*

Context config>subscr-mgmt>mld-policy>static>group

Description This command adds or removes a static multicast source.

The **no** form of this command reverts to the default.

Parameters *grp-ipv6-address* — Specifies the IPv6 address.

Values <grp-ipv6-address>: ipv6-address - x:x:x:x:x:x:x:x (eight 16-bit pieces) x:x:x:x:x:x:d.d.d.d
x - [0 to FFFF]H
d - [0 to 255]D
- multicast group IPv6 address

starg

Syntax [**no**] **starg**

Context config>subscr-mgmt>mld-policy>static>group

Description This command adds a static (*,G) entry. This command can only be enabled if no existing source addresses for this group are specified.

The **no** form of this command removes the starg entry from the configuration.

version

Syntax **version** *version*
no version

Context config>subscr-mgmt>mld-policy#

Description This command configures the MLD version.

The **no** form of this command reverts to the default.

Default version 2

Parameters *version* — Specifies the MLD version.

Values 1, 2

9.25.2.1.15 Multi-Chassis Redundancy Commands

redundancy

Syntax **redundancy**

Context config

Description This command allows the user to perform redundancy operations.

Parameters **force-switchover** — Forces a switchover to the standby CPM card.

Values
now Specifies a keyword to switch to the standby CPM)

Note: Switching to the standby displays the following message.

```
WARNING: Configuration and/or Boot options may have changed since
the last save.
Are you sure you want to switchover (y/n)?
```

synchronize

Syntax **synchronize** {**boot-env** | **config**}

Context config>redundancy

Description This command performs a synchronization of the standby CPM images and/or config files to the active CPM. Either the **boot-env** or **config** parameter must be specified.

In the **config>redundancy** context, this command performs an automatically triggered standby CPM synchronization.

When the standby CPM takes over operation following a failure or reset of the active CPM, it is important to ensure that the active and standby CPM have identical operational parameters. This includes the saved configuration, CPM and IOM images.

The active CPM ensures that the active configuration is maintained on the standby CPM. However, to ensure smooth operation under all circumstances, runtime images and system initialization configurations must also be automatically synchronized between the active and standby CPM.

If synchronization fails, alarms and log messages that indicate the type of error that caused the failure of the synchronization operation are generated. When the error condition ceases to exist, the alarm is cleared.

Only files stored on the router are synchronized. If a configuration file or image is stored in a location other than on a local compact flash, the file is not synchronized (for example, storing a configuration file on an FTP server).

Parameters **boot-env** — Synchronizes all files required for the boot process (loader, BOF, images, and configuration files.

config — Synchronizes only the primary, secondary, and tertiary configuration files.

multi-chassis

Syntax **multi-chassis**

Context config>redundancy

Description This command enables the context to configure multi-chassis parameters.

peer

Syntax **peer** *ip-address* [**create**]
no peer *ip-address*

Context config>redundancy>multi-chassis

Description This command creates a new multi-chassis redundancy peer or enters the context to configure an existing peer.

The peer address can be either an IPv4 address or an IPv6 address. If an IPv6 address is used, the following MCS clients are not supported:

- mc-endpoint
- l2tp

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies an IPv4 or IPv6 unicast address.

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>redundancy>multi-chassis>peer

Description This command configures the authentication key used between this node and the multi-chassis peer. The authentication key can be any combination of letters or numbers.

The **no** form of this command reverts to the default.

Parameters *authentication-key* — Specifies the authentication key. Allowed values are any string up to 20 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

hash-key — Specifies the hash key. The key can be any combination of ASCII characters up to 33 (hash1-key) or 55 (hash2-key) characters (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

hash — Specifies that the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2 — Specifies that the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

mc-ipsec

Syntax **mc-ipsec**

Context config>redundancy>multi-chassis>peer

Description This command enable the context to configure multi-chassis IPSec parameters.

discovery-interval

Syntax **discovery-interval** *interval-1* [**boot** *interval-2*]
no discovery-interval

Context config>redundancy>multi-chassis>peer>mc-ipsec

Description This command specifies the time interval the tunnel group stays in the Discovery state. Interval-1 is used as discovery-interval when a new tunnel group is added to multi-chassis redundancy (mp-ipsec); interval-2 is used as the discovery interval at system boot up. It is optional and when it is not specified, interval-1 is used.

The **no** form of this command reverts to the default.

Default discovery-interval 300 boot 300

Parameters *interval-1/2* — Specifies the interval in seconds.

Values 1 to 1800 seconds

keep-alive-interval

Syntax **keep-alive-interval** *time-interval*
no keep-alive-interval

Context config>redundancy>multi-chassis>peer>mc-ipsec

Description This command specifies the time interval of the mastership election protocol sending the keep-alive packet.

The **no** form of this command reverts to the default.

Default keep-alive-interval 10

Parameters *time-interval* — Specifies the time interval in tenths of a second.

Values 5 to 500

hold-on-neighbor-failure

Syntax **hold-on-neighbor-failure** *multiplier*
no hold-on-neighbor-failure

Context config>redundancy>multi-chassis>peer>mc-ipsec

Description This command specifies the number of keep-alive failures before the peer is considered down.

The **no** form of this command reverts to the default.

Default hold-on-neighbor-failure 3

Parameters *multiplier* — Specifies the number of intervals the standby node will wait for packets from the active node before assuming a redundant-neighbor node failure.

Values 2 to 25

bfd-enable

Syntax [**no**] **bfd-enable**

Context config>redundancy>multi-chassis>peer>mc-ipsec

Description This command enables tracking a central BFD session. If the BFD session goes down, then the system considers the peer down and changes the mc-ipsec status of the configured tunnel-group accordingly.

The BFD session uses the specified loopback interface (in the specified service) address as the source address and uses the specified **dst-ip** as the destination address. Other BFD parameters are configured with the **bfd** command on the specified interface.

The **no** form of this command reverts to the default.

tunnel-group

Syntax **tunnel-group** *group-id* [**create**]
no tunnel-group *group-id*

Context config>redundancy>multi-chassis>peer>mc-ipsec

Description This command enables multi-chassis redundancy for the specified tunnel-group or enters an already configured tunnel-group context. The configured tunnel-group could failover independently.

The **no** form of this command reverts to the default.

Parameters *group-id* — Specifies the tunnel-group ID.

Values 1 to 16

create — Enables multi-chassis redundancy for the specified tunnel-group.

peer-group

Syntax **peer-group** *group-id*
no peer-group

Context config>redundancy>multi-chassis>peer>mc-ipsec>tunnel-group

Description This command specifies the corresponding tunnel-group ID on the peer node. The peer tunnel-group ID does not necessarily equal the local tunnel-group ID.

The **no** form of this command reverts to the default.

Parameters *group-id* — Specifies the tunnel-group ID.

Values 1 to 16

priority

Syntax **priority** *priority*
no priority

Context config>redundancy>multi-chassis>peer>mc-ipsec>tunnel-group

Description This command specifies the local priority of the tunnel group. This is used to elect the master (higher number is the master). If priorities are the same, then the peer with the more active ISA becomes the master. If the priority and the number of active ISAs are the same, then the peer with the higher IP address is the master.

The **no** form of this command reverts to the default.

Parameters *priority* — Specifies the priority of the tunnel-group.

Values 0 to 255

preempt

Syntax [**no**] **preempt**

Context config>redundancy>multi-chassis>peer>mc-ipsec>tunnel-group

Description This command enables the preempt behavior of local node.

The **no** form of this command disables mastership preemption by this tunnel group.

mc-lag

Syntax [**no**] **mc-lag**

Context config>redundancy>multi-chassis>peer>mc-lag

Description This command enables the context to configure multi-chassis LAG operations and related parameters.

The **no** form of this command administratively disables multi-chassis LAG. The **mc-lag** command can only be issued only when MC LAG is shutdown.

hold-on-neighbor-failure

Syntax **hold-on-neighbor-failure** *multiplier*
no hold-on-neighbor-failure

Context config>redundancy>multi-chassis>peer>mc-lag

Description This command specifies the interval that the standby node will wait for packets from the active node before assuming a redundant-neighbor node failure. This delay in switch-over operation is required to accommodate different factors influencing node failure detection rate, such as IGP convergence, or HA switch-over times and to prevent the standby node to take action prematurely.

The **no** form of this command reverts to the default.

Default hold-on-neighbor-failure 3

Parameters *multiplier* — Specifies the time interval that the standby node waits for packets from the active node before assuming a redundant-neighbor node failure.

Values 2 to 25

keep-alive-interval

Syntax **keep-alive-interval** *interval*
no keep-alive-interval

Context config>redundancy>multi-chassis>peer>mc-lag

Description This command sets the interval at which keep-alive messages are exchanged between two systems participating in MC-LAG. These keep-alive messages are used to determine remote-node failure and the interval is set in deci-seconds.

The **no** form of this command reverts to the default.

Default keep-alive-interval 10

Parameters *interval* — Specifies the time interval expressed in deci-seconds.

Values 5 to 500

lag

Syntax **lag** *lag-id* **lacp-key** *admin-key* **system-id** *system-id* [**remote-lag** *lag-id*] **system-priority** *system-priority*
no lag *lag-id*

Context config>redundancy>multi-chassis>peer>mc-lag

Description This command defines a LAG which is forming a redundant-pair for MC-LAG with a LAG configured on the given peer. The same LAG group can be defined only in the scope of 1 peer.

The same **lacp-key**, **system-id**, and **system-priority** must be configured on both nodes of the redundant pair in order to MC-LAG to become operational. In order MC-LAG to become operational, all parameters (**lacp-key**, **system-id**, **system-priority**) must be configured the same on both nodes of the same redundant pair.

The partner system (the system connected to all links forming MC-LAG) will consider all ports using the same **lacp-key**, **system-id**, **system-priority** as the part of the same LAG. In order to achieve this in MC operation, both redundant-pair nodes have to be configured with the same values. In case of the mismatch, MC-LAG is kept in oper-down status.

The **no** form of this command reverts to the default.

Parameters *lag-id* — The LAG identifier, expressed as a decimal integer. Specifying the *lag-id* allows the mismatch between lag-id on redundant-pair. If no **lag-id** is specified it is assumed that neighbor system uses the same *lag-id* as a part of the given MC-LAG. If no matching MC-LAG group can be found between neighbor systems, the individual LAGs will operate normally (no MC-LAG operation is established).

Values 1 to 800

admin-key — Specifies a 16 bit key that needs to be configured in the same manner on both sides of the MC-LAG in order for the MC-LAG to come up.

Values 1 to 65535

system-id — Specifies a 6 byte value expressed in the same notation as MAC address.

Values xx:xx:xx:xx:xx:xx - xx [00 to FF]

lag-id — Specifies the LAG ID on the remote system.

Values 1 to 200 (applies to the 7450 ESS)
1 to 800 (applies to the 7750 SR)

system-priority — Specifies the system priority to be used in the context of the MC-LAG. The partner system will consider all ports using the same **lacp-key**, **system-id**, and **system-priority** as part of the same LAG.

Values 1 to 65535

source-address

Syntax **source-address** *ip-address*
no source-address

Context config>redundancy>multi-chassis>peer

Description This command specifies the source address used to communicate with the multi-chassis peer.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the source address used to communicate with the multi-chassis peer.

sync

Syntax [**no**] **sync**

Context config>redundancy>multi-chassis>peer

Description This command enables the context to configure synchronization parameters.

The **no** form of this command reverts to the default.

ipsec

Syntax [**no**] **ipsec**

Context config>redundancy>multi-chassis>peer>sync

Description This command enables multi-chassis synchronization of IPSec states on system level.

The **no** form of this command reverts to the default.

mc-ring

Syntax **mc-ring**

Context config>redundancy>multi-chassis>peer

Description This command enables the context to configure the multi-chassis ring parameters.

The **no** form of this command reverts to the default.

Default mc-ring

ring

Syntax **ring** *sync-tag* [**create**]
no ring sync-tag

Context config>redundancy>mc>peer>mcr

Description This command configures a multi-chassis ring.

The **no** form of this command removes the sync-tag from the configuration.

l3-ring

Syntax **l3-ring** *name* [**create**]
no l3-ring name

Context config>redundancy>mc>peer>mcr

Description This command configures a Layer 3 multi-chassis ring.

The **no** form of this command reverts to the default.

in-band-control-path

Syntax **in-band-control-path**

Context config>redundancy>mc>peer>mcr>ring
config>redundancy>mc>peer>mc>l3-ring

Description This command enables the context to configure control path parameters.

The **no** form of this command reverts to the default.

debounce

Syntax [**no**] **debounce**

Context config>redundancy>mc>peer>mcr>ring>in-band-control-path
config>redundancy>mc>peer>mc>l3-ring>in-band-control-path

Description This command enables the inband control path debouncing. The **no** form of this command disables inband control path debouncing.

The **no** form of this command reverts to the default.

Default debounce

dst-ip

Syntax **dst-ip** *ip-address*
no dst-ip

Context config>redundancy>mc>peer>mcr>ring>in-band-control-path
config>redundancy>mc>peer>mc>l3-ring>in-band-control-path

Description This command specifies the destination IP address used in the inband control connection.

If the destination IP address is not configured, the ring cannot become operational.

Parameters *ip-address* — The destination IP address.

interface

Syntax **interface** *ip-int-name*
no interface

Context config>redundancy>mc>peer>mcr>ring>in-band-control-path
config>redundancy>mc>peer>mc>l3-ring>in-band-control-path

Description This command specifies the name of the IP interface used for the inband control connection.

If an interface name is not configured, the ring cannot become operational.

The **no** form of this command reverts to the default.

Parameters *ip-int-name* — Specifies an interface name up to 32 characters.

max-debounce-time

Syntax **max-debounce-time** *max-debounce-time*
no max-debounce-time

Context config>redundancy>mc>peer>mcr>ring>in-band-control-path
config>redundancy>mc>peer>mc>l3-ring>in-band-control-path

Description This command configures the inband control path maximum debounce time.

The **no** form of this command reverts to the default.

Default max-debounce-time 10

Parameters *max-debounce-time* — Specifies the maximum debounce time on the transition of the operational state of the inband control connection.

Values 5 to 200 seconds

service-id

Syntax **service-id** *service-id*
no service-id

Context config>redundancy>mc>peer>mcr>ring>ibc

config>redundancy>mc>peer>mcr>l3-ring>ibc

Description This command configures the service ID of the SAP used for the ring node connectivity verification of this ring node. This command supersedes the configuration of a service name.

The no form of this command removes the service-name from the IBC configuration.

Parameters *service-id* — Specifies a service ID or an existing service name.

Values 1 to 214748364 - Only supported in 'classic' configuration-mode (**configure>system>management-interface>configuration-mode classic**)

service-name: A string up to 64 characters

service-name

Syntax **service-name** *service-name*
no service-name

Context config>redundancy>mc>peer>mcr>ring>ibc
config>redundancy>mc>peer>mcr>node>cv
config>redundancy>mc>peer>mc>l3-ring>ibc
config>redundancy>mc>peer>mc>l3-ring>node>cv

Description This command specifies the service name of the SAP used for ring node connectivity verification of this ring node. This command supersedes the configuration of a service ID.

The **no** form of this command removes the service-name from the IBC/CV configuration.

Parameters *service-name* — Specifies the service name up to 64 characters.

path-b

Syntax [**no**] **path-b**

Context config>redundancy>mc>peer>mcr>ring

Description This command specifies the set of upper-VLAN IDs associated with the SAPs that belong to path B with respect to load-sharing. All other SAPs belong to path A.

Default If not specified, the default is an empty set.

range

Syntax [**no**] **range** *vlan-range*

Context config>redundancy>mc>peer>mcr>ring>path-b

config>redundancy>mc>peer>mcr>ring>path-excl

Description This command specifies the set of VLAN IDs associated with the SAPs that are controlled by the remote peer. It is a bitmap that associates bit i with VLAN ID i, with i in [0 to 4094]. Setting the value to the empty string is equivalent to setting it to 512 zeros.

The **no** form of this command reverts to the default.

ring-node

Syntax **ring-node** *ring-node-name* [**create**]
no ring-node *ring-node-name*

Context config>redundancy>mc>peer>mcr>ring

Description This command specifies the unique name of a multi-chassis ring access node.

The **no** form of this command reverts to the default.

Parameters *ring-node-name* — Specifies the ring-node name up to 32 characters.

path-excl

Syntax [**no**] **path-excl**

Context config>redundancy>mc>peer>mcr>ring

Description This command specifies the set of upper-VLAN IDs associated with the SAPs that are to be excluded from control by the multi-chassis ring.

Default If not specified, the default is an empty set.

connectivity-verify

Syntax **connectivity-verify**

Context config>redundancy>mc>peer>mcr>ring
config>redundancy>mc>peer>mc>l3-ring

Description This command enables the context to configure a node connectivity check.

interval

Syntax **interval** *interval*

Context config>redundancy>mc>peer>mcr>ring>cv

config>redundancy>mc>peer>mc>l3-ring>cv

Description This command specifies the polling interval of the ring-node connectivity verification of this ring node.

The **no** form of this command reverts to the default.

Default interval 5

Parameters *interval* — Specifies the polling interval of the ring-node connectivity verification of this ring node.

Values 1 to 6000

service-id

Syntax **service-id** *service-id*
no service-id

Context config>redundancy>mc>peer>mcr>l3ring>in-band-control-path
config>redundancy>mc>peer>mcr>l3ring>node>connectivity-verify
config>redundancy>mc>peer>mcr>ring>in-band-control-path
config>redundancy>mc>peer>mcr>ring>node>connectivity-verify

Description This command specifies the service ID of the SAP used for ring-node connectivity verification of this ring node.

The **no** form of this command reverts to the default.

Parameters *service-id* — Specifies the service ID.

Values service-id: 1 to 214748364

src-ip

Syntax **src-ip** *ip-address*
no src-ip

Context config>redundancy>mc>peer>mcr>ring-node>connect-verify
config>redundancy>mc>peer>mc>l3-ring>connectivity-verify

Description This command specifies the source IP address used in ring-node connectivity verification of this ring node.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the source IP address used in ring-node connectivity verification of this ring node.

service-name

Syntax **service-name** *service-name*
no service-name

Context config>redundancy>mc>peer>mcr>l3ring>in-band-control-path
config>redundancy>mc>peer>mcr>l3ring>node>connectivity-verify
config>redundancy>mc>peer>mcr>ring>in-band-control-path
config>redundancy>mc>peer>mcr>ring>node>connectivity-verify

Description This command specifies the service name of the SAP used for ring-node connectivity verification of this ring node.

The **no** form of this command reverts to the default.

Parameters *service-name* — Specifies the service name up to 64 characters.

src-mac

Syntax **src-mac** *ieee-address*
no src-mac

Context config>redundancy>mc>peer>mcr>ring-node>connect-verify
config>redundancy>mc>peer>mc>l3-ring>connectivity-verify

Description This command specifies the source MAC address used for the Ring-Node Connectivity Verification of this ring node.

If all zeros are specified, then the MAC address of the system management processor (CPM) is used.

The **no** form of this command reverts to the default.

Parameters *ieee-address* — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

vlan

Syntax **vlan** [*vlan-encap*]

Context config>redundancy>mc>peer>mcr>ring-node>connect-verify
config>redundancy>mc>peer>mcr>l3ring>node>cv

Description This command specifies the VLAN tag of the SAP used for ring-node connectivity verification of this ring node. It is only meaningful if the value of is not zero.

The **no** form of this command reverts to the default.

Parameters *vlan-encap* — Specifies the node cc VLAN IP.

srrp-instance

Syntax [**no**] **srrp-instance** *srrp-id*

Context config>redundancy>mc>peer>mc>l3-ring

Description This command configures an SRRP instance for Layer 3 ring.

The **no** form of this command reverts to the default.

Parameters *srrp-id* — Specifies the SRRP ID of this SRRP instance.

Values 1 to 4294967295

9.25.2.1.16 PIM Policy Commands

pim-policy

Syntax **pim-policy** *pim-policy-name* [**create**]
no pim-policy *pim-policy-name*

Context config>subscr-mgmt

Description This command creates a PIM policy or enters the context to configure a PIM policy.

The **no** form of this command deletes the specified PIM policy.

Parameters *pim-policy-name* — Specifies the PIM policy name up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

create — Specifies the keyword used to create the PIM policy. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

pim-policy

Syntax **pim-policy** *policy-name*
no pim-policy *policy-name*

Context config>subscr-mgmt>sub-prof

Description This command adds an existing PIM policy to this subscriber profile.

The **no** form of this command removes the specified PIM policy from this subscriber profile.

Parameters *policy-name* — Specifies the name of the PIM policy name, up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

9.25.2.1.17 RADIUS Accounting Policy Commands

radius-accounting-policy

Syntax **radius-accounting-policy** *name* [**create**]
no radius-accounting-policy

Context config>subscr-mgmt

Description This command specifies a subscriber RADIUS based accounting policy.

The **no** form of this command reverts to the default.

Parameters *name* — Specifies the name of the policy. The string is case sensitive and limited to 32 ASCII 7-bit printable characters with no spaces.

custom-record

Syntax [**no**] **custom-record**

Context config>subscr-mgmt>acct-plcy

Description This command enables the context to configure the layout and setting for a custom accounting record associated with this accounting policy.

The **no** form of this command reverts the configured values to the defaults.

override-counter

Syntax **override-counter** *override-counter-id* [**create**]
no override-counter *override-counter-id*

Context config>log>acct-policy>cr

Description This command enables the context to configure Application Assurance override counter parameters.

The **no** form of this command removes the ID from the configuration.

Parameters *override-counter-id* — Specifies the override counter ID.

Values 1 to 8

e-counters

Syntax **e-counters** [**all**]
no e-counters

Context config>log>acct-policy>cr>override-cntr
config>log>acct-policy>cr>queue
config>log>acct-policy>cr>ref-override-cntr
config>log>acct-policy>cr>ref-queue

Description This command enables the context to configure egress counter parameters for this custom record.

The **no** form of this command reverts to the default.

Parameters **all** — Includes all counters.

i-counters

Syntax **i-counters** [**all**]
no i-counters

Context config>log>acct-policy>cr>override-cntr
config>log>acct-policy>cr>ref-override-cntr
config>log>acct-policy>cr>ref-queue

Description This command configures ingress counter parameters for this custom record.

The **no** form of this command reverts to the default.

Parameters **all** — Includes all counters.

queue

Syntax **queue** *queue-id* [**create**]
no queue *queue-id*

Context config>log>acct-policy>cr

Description This command specifies the queue-id for which counters is collected in this custom record. The counters that is collected are defined in egress and ingress counters.

The **no** form of this command reverts to the default value

Parameters *queue-id* — Specifies the queue-id for which counters is collected in this custom record.

in-profile-octets-discarded-count

Syntax [**no**] **in-profile-octets-discarded-count**

Context config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count

Description This command includes the in-profile octets discarded count.

For queues with **stat-mode v4-v6**, this command includes the IPv4 octets discarded count instead.

The **no** form of this command excludes the in-profile octets discarded count.

in-profile-octets-forwarded-count

Syntax [**no**] **in-profile-octets-forwarded-count**

Context config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count

Description This command includes the in-profile octets forwarded count. For queues with **stat-mode v4-v6**, this command includes the IPv4 octets forwarded count instead.

The **no** form of this command excludes the in-profile octets forwarded count.

in-profile-packets-discarded-count

Syntax [**no**] **in-profile-packets-discarded-count**

Context config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count

Description This command includes the in-profile packets discarded count.

For queues with **stat-mode v4-v6**, this command includes the IPv4 packets discarded count instead.

The **no** form of this command excludes the in-profile packets discarded count.

in-profile-packets-forwarded-count

Syntax [**no**] **in-profile-packets-forwarded-count**

Context config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count

Description This command includes the in-profile packets forwarded count.

For queues with **stat-mode v4-v6**, this command includes the IPv4 packets forwarded count instead.

The **no** form of this command excludes the in-profile packets forwarded count.

out-profile-octets-discarded-count

Syntax [**no**] **out-profile-octets-discarded-count**

Context config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count

Description This command includes the out of profile packets discarded count.

For queues with **stat-mode v4-v6**, this command includes the IPv6 octets discarded count instead.

The **no** form of this command excludes the out of profile packets discarded count.

out-profile-octets-forwarded-count

Syntax [**no**] **out-profile-octets-forwarded-count**

Context config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count

Description This command includes the out of profile octets forwarded count.

For queues with **stat-mode v4-v6**, this command includes the IPv6 octets forwarded count instead.

The **no** form of this command excludes the out of profile octets forwarded count.

out-profile-packets-discarded-count

Syntax [**no**] **out-profile-packets-discarded-count**

Context config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count

Description This command includes the out of profile packets discarded count.

For queues with **stat-mode v4-v6**, this command includes the IPv6 packets discarded count instead.

The **no** form of this command excludes the out of profile packets discarded count.

out-profile-packets-forwarded-count

Syntax [**no**] **out-profile-packets-forwarded-count**

Context config>log>acct-policy>cr>oc>e-count
config>log>acct-policy>cr>roc>e-count
config>log>acct-policy>cr>queue>e-count
config>log>acct-policy>cr>ref-queue>e-count

Description This command includes the out of profile packets forwarded count.

For queues with **stat-mode v4-v6**, this command includes the IPv6 packets forwarded count instead.

The **no** form of this command excludes the out of profile packets forwarded count.

all-octets-offered-count

Syntax [**no**] **all-octets-offered-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes all octets offered in the count.

The **no** form of this command excludes the octets offered in the count.

all-packets-offered-count

Syntax [**no**] **all-packets-offered-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes all packets offered in the count.

The **no** form of this command excludes the packets offered in the count.

high-octets-discarded-count

Syntax [**no**] **high-octets-discarded-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the high octets discarded count.

For queues with **stat-mode v4-v6**, this command includes the IPv4 octets discarded count instead.

The **no** form of this command excludes the high octets discarded count.

high-octets-offered-count

Syntax [**no**] **high-octets-offered-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the high octets offered count.

The **no** form of this command excludes the high octets offered count.

high-packets-discarded-count

Syntax [**no**] **high-packets-discarded-count**

Context config>log>acct-policy>cr>oc>i-count

config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the high packets discarded count.

For queues with **stat-mode v4-v6**, this command includes the IPv4 packets discarded count instead.

The **no** form of this command excludes the high packets discarded count.

high-packets-offered-count

Syntax [**no**] **high-packets-offered-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the high packets offered count.

The **no** form of this command excludes the high packets offered count.

in-profile-octets-forwarded-count

Syntax [**no**] **in-profile-octets-forwarded-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the in profile octets forwarded count.

For queues with **stat-mode v4-v6**, this command includes the IPv4 octets forwarded count instead.

The **no** form of this command excludes the in profile octets forwarded count.

in-profile-packets-forwarded-count

Syntax [**no**] **in-profile-packets-forwarded-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count

config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the in profile packets forwarded count.

For queues with **stat-mode v4-v6**, this command includes IPv4 packets forwarded count instead.

The **no** form of this command excludes the in profile packets forwarded count.

low-octets-discarded-count

Syntax [**no**] **low-octets-discarded-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the low octets discarded count.

For queues with **stat-mode v4-v6**, this command includes the IPv6 octets discarded count instead.

The **no** form of this command excludes the low octets discarded count.

low-packets-discarded-count

Syntax [**no**] **low-packets-discarded-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the low packets discarded count.

For queues with **stat-mode v4-v6**, this command includes the IPv6 packets discarded count instead.

The **no** form of this command excludes the low packets discarded count.

low-octets-offered-count

Syntax [**no**] **low-octets-offered-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count

config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the low octets discarded count.

The **no** form of this command excludes the low octets discarded count.

low-packets-offered-count

Syntax [**no**] **low-packets-offered-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the low packets discarded count.

The **no** form of this command excludes the low packets discarded count.

out-profile-octets-forwarded-count

Syntax [**no**] **out-profile-octets-forwarded-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the out of profile octets forwarded count.

For queues with **stat-mode v4-v6**, this command includes the IPv6 octets forwarded count instead.

The **no** form of this command excludes the out of profile octets forwarded count.

out-profile-packets-forwarded-count

Syntax [**no**] **out-profile-packets-forwarded-count**

Context config>log>acct-policy>cr>oc>i-count
config>log>acct-policy>cr>roc>i-count
config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the out of profile packets forwarded count.

For queues with **stat-mode v4-v6**, this command includes the IPv6 packets forwarded count instead.

The **no** form of this command excludes the out of profile packets forwarded count.

uncoloured-octets-offered-count

Syntax [**no**] **uncoloured-packets-offered-count**

Context config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the uncoloured octets offered in the count.

The **no** form of this command excludes the uncoloured octets offered in the count.

uncoloured-packets-offered-count

Syntax [**no**] **uncoloured-packets-offered-count**

Context config>log>acct-policy>cr>queue>i-count
config>log>acct-policy>cr>ref-queue>i-count

Description This command includes the uncoloured packets offered count.

The **no** form of this command excludes the uncoloured packets offered count.

ref-aa-specific-counter

Syntax [**no**] **ref-aa-specific-counter any**

Context config>log>acct-policy>cr

Description This command includes a reference aa counter.

The **no** form of this command reverts to the default value.

ref-override-counter

Syntax **ref-override-counter** *ref-override-counter-id*
ref-override-counter all
no ref-override-counter

Context config>log>acct-policy>cr

Description This command configures a reference override counter.

The **no** form of this command reverts to the default value.

Parameters *ref-override-counter-id* — Specifies an override counter ID.

Values 1 to 8

all — Specifies to include all override counters.

ref-queue

Syntax **ref-queue** *queue-id*
ref-queue all
no ref-queue

Context config>log>acct-policy>cr

Description This command configures a reference queue.

The **no** form of this command reverts to the default value.

significant-change

Syntax **significant-change** *delta*
no significant-change

Context config>log>acct-policy>cr

Description This command configures the significant change required to generate the record.

Parameters *delta* — Specifies the delta change (significant change) that is required for the custom record to be written to the XML file.

Values 0 to 4294967295

delay-start-time

Syntax [**no**] **delay-start-time**

Context config>subscr-mgmt>acct-plcy

Description It is recommended to only use this command when **session-accounting** is used. By default, a dual stack subscriber generates a RADIUS accounting interim message for each new host update (IPv4, IPv6 WAN, and IPv6 PD). This command delays the trigger of a RADIUS accounting start message and allows all hosts to connect first. When the delay timer expires, a single RADIUS accounting start message containing all the host currently connected to the BNG is sent to the server. Subsequent host connections will trigger interim-updates if host-update is enabled on session-accounting. For all other accounting modes, this command will delay the trigger of an accounting start when a host connects.

The **no** form of this command reverts to the default.

9.25.2.1.18 RADIUS Route Download Commands

route-downloader

Syntax **route-downloader** *name* [**create**]
no route-downloader *name*

Context config>aaa

Description This command enables the context to configure a route-downloader instance. The route-downloader is a process that uses radius access-request messages to a particular server. The server returns either an access-accept or access-deny message. Access-accept messages also contain the prefixes (in the form of static blackhole routes in various formats). Only a single route-downloader object can be created.

The **no** form of this command removes the name from the configuration. The object must be shutdown prior to deletion. No prefix is needed to delete an existing route-download object.

Parameters *name* — Specifies the name of this RADIUS route downloader.

create — This keyword is mandatory while creating an instance of the route-download object.

base-user-name

Syntax **base-user-name** *user-name*
no base-user-name

Context config>aaa>route-downloader

Description This command sets the prefix for the user name that shall be used as access requests. The actual name used is a concatenation of this string, the “-” (dash) character and a monotonically increasing integer.

The **no** form of this command removes the *user-name* from the configuration.

Default The system’s configured name (system name).

Parameters *user-name* — Specifies the prefix of the username that is used in the RADIUS access requests. The username used in the RADIUS access requests is a concatenation of this string, the dash character and an increasing integer.

default-metric

Syntax **default-metric** *metric*
no default-metric

Context config>aaa>route-downloader

Description This command sets the default metric that routes imported by the RTM will acquire.

The **no** form of this command removes the metric

Default default-metric 2

Parameters *metric* — Specifies the default metric of the routes imported.

Values 0 to 254

default-tag

Syntax **default-tag** *tag*
no default-tag

Context config>aaa>route-downloader

Description This command sets the default tag that routes processed by the AAA route downloader will take.

Note: Any route received with a specific tag retains the specific tag. The tag value is passed to the Route Table Manager and is available as match condition on the export statement of other routing protocols.

The **no** form of this command reverts to the default.

Default default-tag 0

Parameters *tag* — Specifies the default tag of the routes imported.

Values 0 to 4294967295

download-interval

Syntax **download-interval** *minutes*
no download-interval

Context config>aaa>route-downloader

Description This command sets the time interval, in minutes, that the system waits for between two consecutive runs of the route-download process. The time is counted from the start-time of the run, thus, if an route-download process is still ongoing by the time the timer expires, the process will restart from count=1.

The **no** form of this command reverts to the default value.

Default download-interval 720

Parameters *minutes* — Specifies the time interval, in minutes, between the start of the last route downloader run and the start of the next route downloader run.

Values 1 to 1440

max-routes

Syntax **max-routes** *routes*
no max-routes

Context config>aaa>route-downloader

Description This command determines the upper limits for total number of routes to be received and accepted by the system. The total number is inclusive of both IPv4 and IPv6 addresses and no differentiation is needed across protocols. It includes the sum of both. Once this limit is reached, the download process stops sending new access-requests until the next download-interval expires.

The **no** form of this command reverts to the default value.

Default max-routes 200000

Parameters *routes* — Specifies the maximum number of the routes imported.

Values 1 to 200000

password

Syntax **password** *password* [**hash** | **hash2**]
no password

Context config>aaa>route-downloader

Description This command specifies the password that is used in the RADIUS access requests. It shall be specified as a string of up to 32 characters.

The **no** form of this command resets the password to its default of **ALU** and is stored using hash/hash2 encryption.

Default password ALU

Parameters *password* — Specifies a password string up to 32 characters.

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

radius-server-policy

Syntax **radius-server-policy** *policy-name*
no radius-server-policy

Context config>aaa>route-downloader

Description This command references an existing radius-server-policy (available under the **config>aaa** context). The server (or servers) referenced by the policy is used as the targets for the access-request message.

The **no** form of this command removes the policy name from the route-downloader configuration.

Parameters *policy-name* — Specifies the RADIUS server policy.

retry-interval

Syntax **retry-interval min** *minimum* **max** *maximum*
no retry-interval

Context config>aaa>route-downloader

Description This command sets the duration, in minutes, of the retry interval. The retry interval is the interval meant for the system to retry sending an Access Request message after the previous one was unanswered (not with an access reject but rather just a RADIUS failure or ICMP port unreachable). This timer is actually an exponential backoff timer that starts at **min** and is capped at **max** minutes.

The **no** form of this command reverts to the default values.

Default retry-interval min 10 max 20

Parameters **min** *minimum* — Specifies the duration, in minutes, of the retry interval. This duration grows exponentially after each sequential failure.

Values 1 to 1440

Default 10

max *maximum* — Specifies the maximum duration, in minutes, of the retry interval.

Values 1 to 1440

Default 20

9.25.2.1.19 Redundant Interface Commands

Note: The commands described in this section apply only to the 7750 SR.

redundant-interface

Syntax [**no**] **redundant-interface** *ip-int-name*

Context config>service>ies
config>service>vprn
config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command configures a redundant interface.

The **no** form of this command removes the value from the configuration.

Parameters *ip-int-name* — Specifies the name of the IP interface. Interface names can be from 1 to 32 alphanumeric characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

remote-proxy-arp

Syntax [**no**] **remote-proxy-arp**

Context config>service>ies>if
config>service>ies>sub-if>grp-if
config>service>vprn>if
config>service>vprn>sub-if>grp-if

Description This command enables remote proxy ARP on the interface.

Remote proxy ARP is similar to proxy ARP. It allows the router to answer an ARP request on an interface for a subnet that is not provisioned on that interface. This allows the router to forward to the other subnet on behalf of the requester. To distinguish remote proxy ARP from local proxy ARP, local proxy ARP performs a similar function but only when the requested IP is on the receiving interface.

The **no** form of this command reverts to the default.

address

Syntax **address** {*ip-address/mask* | *ip-address netmask*} [**remote-ip** *ip-address*]
no address

Context config>service>vprn>redundant-if

Description This command assigns an IP address mask or netmask and a remote IP address to the interface.

The **no** form of this command removes the values from the configuration.

Parameters *ip-address/mask* — Assigns an IP address/IP subnet format to the interface.

ip-address netmask — Assigns an IP address netmask to the interface. Specifies a string of 0s and 1s that mask or screen out the network part of an IP address so that only the host computer part of the address remains.

remote-ip *ip-address* — Assigns a remote IP to the interface.

spoke-sdp

Syntax [**no**] **spoke-sdp** *sdp-id*

Context config>service>vprn

Description This command binds a service to an existing Service Distribution Point (SDP).

A spoke SDP is treated like the equivalent of a traditional bridge “port” where flooded traffic received on the spoke SDP is replicated on all other “ports” (other spoke and mesh SDPs or SAPs) and not transmitted on the port it was received.

The SDP has an operational state which determines the operational state of the SDP within the service. For example, if the SDP is administratively or operationally down, the SDP for the service is down.

The SDP must already be defined in the **config>service>sdp** context in order to associate an SDP with a VPRN service. If the **sdp** *sdp-id* is not already configured, an error message is generated. If the *sdp-id* does exist, a binding between that *sdp-id* and the service is created.

SDPs must be explicitly associated and bound to a service. If an SDP is not bound to a service, no far-end 7750 SR devices can participate in the service.

The **no** form of this command removes the SDP binding from the service. The SDP configuration is not affected; only the binding of the SDP to a service. Once removed, no packets are forwarded to the far-end router.

Special Cases **VPRN —** Several SDPs can be bound to a VPRN service. Each SDP must be destined to a different 7750 SR OS. If two *sdp-id* bindings terminate on the same 7750 SR, an error occurs and the second SDP binding is rejected.

Parameters *sdp-id* — Specifies the SDP identifier. Allowed values are integers in the range of 1 and 17407 for existing SDPs.

vc-id — Specifies the virtual circuit identifier.

Values 1 to 4294967295

egress

Syntax **egress**

Context config>service>vprn>red-if>spoke-sdp

Description This command enables the context to configure egress SDP parameters.

ingress

Syntax **ingress**

Context config>service>vprn>red-if>spoke-sdp

Description This command enables the context to configure ingress SDP parameters.

vc-label

Syntax **vc-label** *egress-vc-label*
no vc-label [*egress-vc-label*]

Context config>service>vprn>red-if>spoke-sdp>egress

Description This command configures the egress VC label.

The **no** form of this command removes the value from the configuration.

Parameters *vc-label* — Specifies the VC egress value which indicates a specific connection.

Values 16 to 1048575

vc-label

Syntax **vc-label** *ingress-vc-label*
no vc-label [*ingress-vc-label*]

Context config>service>vprn>red-if>spoke-sdp>ingress

Description This command configures the ingress VC label.

The **no** form of this command removes the value from the configuration.

Parameters *vc-label* — Specifies a VC ingress value which indicates a specific connection.

Values 2048 to 18431

filter

Syntax **filter** {**ip** *ip-filter-id*}
no filter

Context config>service>vprn>red-if>spoke-sdp>ingress
config>service>vprn>red-if>spoke-sdp>egress

Description This command associates an IP filter policy with an ingress or egress Service Access Point (SAP) or IP interface. An IP filter policy can be associated with spoke SDPs.

Filter policies control the forwarding and dropping of packets based on IP or MAC matching criteria.

The filter command is used to associate a filter policy with a specified ip-filter-id with an ingress or egress SAP. The *ip-filter-id* must already be defined before the filter command is executed. If the filter policy does not exist, the operation fail sand an error message returned.

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to IP match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP or IP interface. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use scope command within the filter definition to change the scope to local or global. The default scope of a filter is local.

Parameters *ip-filter-id* — Specifies an existing IP filter policy.

Values 1 to 65535

9.25.2.1.20 RIP Commands

rip-policy

Syntax **rip-policy** *policy-name* [**create**]
no rip-policy *policy-name*

Context config>subscr-mgmt

Description This command creates a RIP policy. This policy is applied to a subscriber IPv4 host to enable the BNG to learn RIP routes from the host. RIP routes are never sent to the hosts.

Parameters *policy-name* — Specifies the RIP policy name up to 32 characters.

create — Keyword required to create the configuration context.

neighbor

Syntax [**no**] **neighbor** *ip-int-name*

Context config>router>rip>group
config>service>vprn>rip>group

Description This command creates a context for configuring a RIP neighbor interface. By default, group interfaces are not activated with RIP, unless explicitly configured. The BNG only learns RIP routes from IPv4 host on the group interface. The RIP neighbor group interface defaults to **none**. The send operation is unchangeable for group-interface.

The **no** form of this command deletes the RIP interface configuration for this group interface. The shutdown command in the **config>router>rip>group group-name>neighbor** context can be used to disable an interface without removing the configuration for the interface.

Parameters *ip-int-name* — Specifies the group interface name. Interface names must be unique within the group of defined group interfaces within **config>service>vprn/ies sub-if> grp-if** commands. An interface name cannot be in the form of an IP address. Interface names can be any string, up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes. If the group interface name does not exist, an error message is returned.

authentication-key

Syntax **authentication-key** [*authentication-key* | *hash-key*] [**hash** | **hash2**]
no authentication-key

Context config>subscr-mgmt>rip-policy

Description This command configures the BGP authentication key.

Authentication is performed between neighboring routers before setting up the BGP session by verifying the password. Authentication is performed using the MD-5 message-based digest. The authentication key can be any combination of letters or numbers from 1 to 16.

The **no** form of this command removes the authentication password from the configuration and effectively disables authentication.

Default Authentication is disabled and the authentication password is empty.

Parameters *authentication-key* — Specifies the authentication key. The key can be any combination of ASCII characters up to 255 characters (unencrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

hash-key — Specifies the hash key. The key can be any combination of ASCII characters up to 342 characters (encrypted). If spaces are used in the string, enclose the entire string in quotation marks (“ ”).

This is useful when a user must configure the parameter, but, for security purposes, the actual unencrypted key value is not provided.

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

authentication-type

Syntax **authentication-type** {**none** | **password** | **message-digest** | **message-digest-20**}
no authentication-type

Context config>sub-mgmg>rip-policy>

Description This command sets the type of authentication to be used between RIP neighbors. The type and password must match exactly for the RIP message to be considered authentic and processed.

The **no** form of this command removes the authentication type from the configuration and effectively disables authentication.

Parameters **none** — Disables authentication at a given level (global, group, neighbor). If the command does not exist in the configuration, the parameter is inherited.

password — Specifies enable simple password (plain text) authentication. If authentication is enabled and no authentication type is specified in the command, simple password authentication is enabled.

message-digest — Configures 16 byte message digest for MD5 authentication. If this option is configured, then at least one message-digest-key must be configured.

message-digest-20 — Configures 20 byte message digest for MD5 authentication in accordance with RFC 2082, RIP-2 MD5 Authentication. If this option is configured, then at least one message-digest-key must be configured.

retail-svc-id

Syntax **retail-svc-id** *service-id*
retail-svc-id

Context config>service>ies>sub-if>grp-if>sap>static-host
config>service>vprn>sub-if>grp-if>sap>static-host

Description This command specifies the service id of the retailer IES/VPRN service to which the static IPv6 host belongs. A corresponding retailer subscriber interface must exist in the specified service.

The **no** form of this command reverts to the default.

Parameters *service-id* — Specifies the retailer service ID or retailer service name.

Values *service-id*: 1 to 2148007978
svc-name: A string up to 64 characters

sla-profile

Syntax **sla-profile** *sla-profile-name*
no sla-profile

Context config>service>ies>if>sap>static-host
config>service>vpls>sap>static-host
config>service>ies>sub-if>grp-if>sap>static-host
config>service>vprn>if>sap>static-host
config>service>vprn>sub-if>grp-if>sap>static-host

Description This command specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

The **no** form of this command reverts to the default.

Parameters *sla-profile-name* — Specifies the SLA profile name.

sub-profile

Syntax **sub-profile** *sub-profile-name*
no sub-profile

Context config>service>ies>if>sap>static-host
config>service>vpls>sap>static-host
config>service>ies>sub-if>grp-if>sap>static-host
config>service>vprn>if>sap>static-host
config>service>vprn>sub-if>grp-if>sap>static-host

Description This command specifies an existing subscriber profile name to be associated with the static subscriber host.

The **no** form of this command reverts to the default.

Parameters *sub-profile-name* — Specifies the sub-profile name.

subscriber

Syntax **subscriber** *sub-ident*
no subscriber

Context config>service>ies>if>sap>static-host
config>service>vpls>sap>static-host
config>service>ies>sub-if>grp-if>sap>static-host
config>service>vprn>if>sap>static-host
config>service>vprn>sub-if>grp-if>sap>static-host

Description This command specifies an existing subscriber identification profile to be associated with the static subscriber host.

Parameters *sub-ident* — Specifies the subscriber identification.

subscriber-sap-id

Syntax [**no**] **subscriber-sap-id**

Context config>service>ies>if>sap>static-host
config>service>ies>sub-if>grp-if>sap>static-host
config>service>vpls>sap>static-host
config>service>vprn>if>sap>static-host
config>service>vprn>sub-if>grp-if>sap>static-host

Description This command enables using the SAP ID as subscriber id.

Parameters **subscriber-sap-id** — Specifies to use the sap-id as the subscriber-id.

rip

Syntax [**no**] **rip**

Context config>service>vprn
config>service>ies

Description This command enables the RIP protocol on the given VPRN IP interface.

The **no** form of this command disables the RIP protocol from the given VPRN IP interface.

group

Syntax [**no**] **group** *group-name*

Context config>service>vprn>rip
config>service>ies>rip

Description This command creates a context for configuring a RIP group of neighbors. RIP groups are a way of logically associating RIP neighbor interfaces to facilitate a common configuration for RIP interfaces.

The **no** form of this command deletes the RIP neighbor interface group. Deleting the group also removes the RIP configuration of all the neighbor interfaces currently assigned to this group.

Parameters *group-name* — The RIP group name. Allowed values are any string, up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

9.25.2.1.21 Router Advertisement Policy Commands

router-advertisement-policy

Syntax **router-advertisement-policy** *name* [**create**]
no router-advertisement-policy *name*

Context config>subscr-mgmt

Description This command creates a router advertisement policy or enters the context to configure a router advertisement policy. The keyword **create** is mandatory when creating a router advertisement policy the first time.

The **no** form of this command deletes the specified router advertisement policy.

Parameters *name* — Specifies the router advertisement policy name up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

create — Specifies the keyword used to create the router advertisement policy. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

router-advertisements

Syntax [**no**] **router-advertisements**

Context config>service>ies>sub-if>ipv6
config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6
config>service>vprn>sub-if>grp-if>ipv6

Description This command enables the context to configure IPv6 router advertisements for this interface.

current-hop-limit

Syntax **current-hop-limit** *limit*
no current-hop-limit

Context config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6>rtr-adv
config>service>vprn>router-advert>if
config>service>vprn>sub-if>ipv6>rtr-adv

Description This command configures the hop limit to be advertised.

The **no** form of this command returns the command to the default setting.

Default current-hop-limit 64

Parameters *limit* — Specifies the default value to be placed in the current hop limit field in router advertisement policies sent.

Values 0 to 255

dns-options

Syntax [**no**] **dns-options**

Context config>subscr-mgmt>rtr-adv-plcy
config>service>ies>sub-if>ipv6>rtr-adv

config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6

Description This command enables the context to configure IPv6 DNS options for SLAAC hosts.

The **no** form of this command returns the command to the default setting.

include-dns

Syntax [**no**] **include-dns**

Context config>subscr-mgmt>rtr-adv-plcy>dns-opt
config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6

Description This command specifies to include the Recursive DNS Server (RDNSS) Option as defined in RFC 6106 in IPv6 router advertisements for DNS name resolution of IPv6 SLAAC hosts.

The **no** form of this command returns the command to the default setting.

rdnss-lifetime

Syntax **rdnss-lifetime** *seconds*
rdnss-lifetime infinite
no rdnss-lifetime

Context config>subscr-mgmt>rtr-adv-plcy>dns-opt
config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6

Description This command configures the maximum time that the RDNSS address may be used for name resolution.

The **no** form of this command returns the command to the default setting.

Default rdnss-lifetime 3600

Parameters *seconds* — Specifies the time, in seconds, that the RDNSS address is valid for this route.

Values 900 to 3600

infinite — Specifies that the RDNSS address can be used permanently.

force-mcast

Syntax **force-mcast** [**ip**] [**mac**]
no force-mcast

Context config>subscr-mgmt>rtr-adv-plcy
config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6

Description This command configures the protocols with forced multicast, either IP or MAC.

The **no** form of this command returns the command to the default setting.

Parameters **ip** — Specifies that IP for multicast will be forced.

mac — Specifies that MAC for multicast will be forced.

managed-configuration

Syntax [**no**] **managed-configuration**

Context config>subscr-mgmt>rtr-adv-plcy
config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6

Description This command sets or resets the managed address configuration flag.

The **no** form of this command returns the command to the default setting.

max-advertisement

Syntax **max-advertisement** *seconds*
no max-advertisement

Context config>subscr-mgmt>rtr-adv-plcy
config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6

Description This command specifies the maximum time allowed between sending unsolicited router advertisements.

The **no** form of this command returns the command to the default setting.

Default max-advertisement 1800

Parameters *seconds* — Specifies the maximum advertisement interval, in seconds.

Values 900 to 1800

min-advertisement

Syntax **max-advertisement** *seconds*
no max-advertisement

Context config>subscr-mgmt>rtr-adv-plcy
config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6

Description This command specifies the minimum time allowed between sending unsolicited router advertisements.

The **no** form of this command returns the command to the default setting.

Default min-advertisement 900

Parameters *seconds* — Specifies the minimum advertisement interval, in seconds.

Values 900 to 1350

mtu

Syntax **mtu** *bytes*
no mtu

Context config>subscr-mgmt>rtr-adv-plcy
config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6

Description This command specifies the value to be placed in link MTU options sent by the router.

The **no** form of this command returns the command to the default setting.

Parameters *bytes* — Sets the advertised MTU value, in bytes, for this interface.

Values 1280 to 9212

other-stateful-configuration

Syntax [**no**] **other-stateful-configuration**

Context config>subscr-mgmt>rtr-adv-plcy
config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6

Description This command sets and resets the other-stateful-configuration flag.

The **no** form of this command returns the command to the default setting.

prefix-options

Syntax [**no**] **prefix-options**

Context config>subscr-mgmt>rtr-adv-plcy
config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6

Description This command enters the configuration of prefix options.

The **no** form of this command reverts to the default.

autonomous

Syntax [**no**] **autonomous**

Context config>service>ies>sub-if>ipv6>rtr-adv>pfx-opt
config>service>ies>sub-if>grp-if>ipv6>rtr-adv>pfx-opt
config>service>vprn>sub-if>ipv6>rtr-adv>pfx-op
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv>pfx-opt

Description This command enables the option that determines whether or not the prefix can be used for stateless address autoconfiguration.

The **no** form of this command disables the option.

on-link

Syntax [**no**] **on-link**

Context config>router>router-advert>if

config>service>ies>sub-if>ipv6>rtr-adv>pfx-opt
config>service>ies>sub-if>grp-if>ipv6>rtr-adv>pfx-opt
config>service>vprn>router-advert>if
config>service>vprn>sub-if>ipv6>rtr-adv>pfx-opt
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv>pfx-opt

Description This command specifies whether the prefix is assigned to an interface on the specified link.

The **no** form of this command reverts to the default.

Default on-link

preferred-lifetime

Syntax **preferred-lifetime** *seconds*
preferred-lifetime infinite
no preferred-lifetime

Context config>service>ies>sub-if>ipv6>rtr-adv>pfx-opt
config>service>ies>sub-if>grp-if>ipv6>rtr-adv>pfx-opt
config>service>vprn>sub-if>ipv6>rtr-adv>pfx-opt
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv>pfx-opt

Description This command specifies the remaining time for this prefix to be preferred, thus time until deprecation.

The **no** form of this command reverts to the default.

Default preferred-lifetime 3600

Parameters *seconds* — Specifies the time, in seconds, for the prefix to remain preferred on this interface in seconds.

Values 0 to 4294967295

infinite — Specifies that the remaining time will never expire.

stateful

Syntax [**no**] **stateful**

Context config>subscr-mgmt>rtr-adv-plcy>pfx-opt

Description This command enables the configuration of RA options for stateful DHCP prefixes used by the subscriber host.

The **no** form of this command reverts to the default.

stateless

Syntax [**no**] **stateless**

Context config>subscr-mgmt>rtr-adv-plcy>pfx-opt

Description This command enables the configuration of RA options for stateless SLAAC prefixes used by the subscriber host.

The **no** form of this command reverts to the default.

auto-lifetimes

Syntax [**no**] **auto-lifetimes**

Context config>subscr-mgmt>rtr-adv-plcy>pfx-opt>stateful

Description This command adjusts the valid and preferred lifetime values of the router advertisement from the DHCP lease of the subscriber. Every router advertisement sent to the subscriber is derived from the DHCP lease in real time. The route advertisement is always sent on a DHCP Renew.

The **no** form of this command reverts to the default.

on-link

Syntax [**no**] **on-link**

Context config>subscr-mgmt>rtr-adv>pfx-opt>stateful
config>subscr-mgmt>rtr-adv>pfx-opt>stateless

Description This command specifies whether the prefix is to be assigned to an interface on the specified link.

The **no** form of this command reverts to the default.

Default on-link

preferred-lifetime

Syntax **preferred-lifetime** *seconds*
preferred-lifetime infinite
no preferred-lifetime

Context config>subscr-mgmt>rtr-adv>pfx-opt>stateful
config>subscr-mgmt>rtr-adv>pfx-opt>stateless

Description This command specifies the remaining time for this prefix to be preferred.

The **no** form of this command reverts to the default.

Default preferred-lifetime 3600

Parameters *seconds* — Specifies the time, in seconds, for the prefix to remain preferred.

Values 0, 900 to 86400

infinite — Specifies that the remaining time never expires.

valid-lifetime

Syntax **preferred-lifetime** *seconds*
preferred-lifetime infinite
no preferred-lifetime

Context config>subscr-mgmt>rtr-adv>pfx-opt>stateful
config>subscr-mgmt>rtr-adv>pfx-opt>stateless

Description This command specifies the time for this prefix to remain valid.

The **no** form of this command reverts to the default.

Default valid-lifetime 86400

Parameters *seconds* — Specifies the time, in seconds, for the prefix to remain preferred.

Values 0, 900 to 86400

infinite — Specifies that the time never expires.

reachable-time

Syntax **reachable-time** *milli-seconds*
no reachable-time

Context config>subscr-mgmt>rtr-adv-plcy

Description This command configures the reachable time for advertisements.

The **no** form of this command returns the command to the default setting.

Default reachable-time 0

Parameters *milli-seconds* — Specifies the time, in milliseconds, for the reachable time.

Values 0 to 3600000

retransmit-time

Syntax **retransmit-time** *milli-seconds*
no retransmit-time

Context config>subscr-mgmt>rtr-adv-plcy

Description This command configures the retransmit time for advertisements.

The **no** form of this command returns the command to the default setting.

Default retransmit-time 0

Parameters *milli-seconds* — Specifies the time, in milliseconds, for the retransmit time.

Values 0 to 1800000

router-lifetime

Syntax **router-lifetime** *seconds*
router-lifetime no-default-router
no router-lifetime

Context config>subscr-mgmt>rtr-adv-plcy

Description This command specifies the router lifetime.

The **no** form of this command returns the command to the default setting.

Default router-lifetime 4500

Parameters *seconds* — Specifies the time, in seconds, for the prefix to remain preferred.

Values 2700 to 9000

no-default-router — Specifies that the router is not to be used as a default router.

router-solicit

Syntax **router-solicit**

Context config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6
config>service>ies>sub-if>ipv6

Description This command enables the context to configure parameters used for router-solicit based authentication.

inactivity-timer

Syntax **inactivity-timer** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no inactivity-timer

Context config>service>ies>sub-if>ipv6>rtr-sol
config>service>ies>sub-if>grp-if>ipv6>rtr-sol
config>service>vprn>sub-if>ipv6>rtr-sol
config>service>ies>sub-if>grp-if>ipv6>rtr-sol

Description This command specifies the time before an inactive host is removed.

The **no** form of this command reverts to the default.

Parameters **infinite** — Specifies that the idle host is never removed.

days — Specifies that the idle host is removed if idle within the number of specified days.

hours — Specifies that the idle host is removed if idle within the number of specified hours.

minutes — Specifies that the idle host is removed if idle within the number of specified minutes.

seconds — Specifies that the idle host is removed if idle within the number of specified seconds.

9.25.2.1.22 SLA Profile Commands

sla-profile

Syntax **sla-profile** *sla-profile-name* [**create**]
no sla-profile *sla-profile-name*

Context config>subscr-mgmt

Description This command configures an SLA profile mapping. Hosts associated with a subscriber are subdivided into Service Level Agreement (SLA) profiles. For each subscriber host an SLA profile can be specified. For a subscriber host, the SLA profile determines:

- The QoS-policies to use
  - The classification
  - The queues
  - The queue mapping
- The IP filters to use

The SLA profile also has the attribute host-limits which limits the total number of hosts (belonging to the same subscriber) on a certain SAP that can be using this SLA profile.

The **no** form of this command reverts to the default.

Parameters *sla-profile-name* — Specifies the name of the SLA profile.

egress

Syntax **egress**

Context config>subscr-mgmt>sla-profile

Description This command enables the context to configure egress parameters for the SLA profile.

ingress

Syntax **ingress**

Context config>subscr-mgmt>sla-profile

Description This command enables the context to configure ingress parameters for the SLA profile.

host-limits

Syntax [**no**] **no host-limits**

Context config>subscr-mgmt>sla-profile

Description This command enables the context to configure host limit parameter for this for this SLA profile.

The **no** form of this command removes the parameters from the configuration.

ipv4-arp

Syntax **ipv4-arp** *max-nr-of-hosts*
no ipv4-arp

Context config>subscr-mgmt>sla-profile>host-limits

Description This command configures the maximum number of IPv4 ARP hosts.

The **no** form of this command removes the number of IPv4 ARP hosts from the SLA profile.

Parameters *max-nr-of-hosts* — Specifies the maximum number of IPv4 ARP hosts.

Note: The operational maximum value may be smaller due to equipped hardware dependencies

Values 0 to 131071

ipv4-dhcp

Syntax **ipv4-dhcp** *max-nr-of-hosts*
no ipv4-dhcp

Context config>subscr-mgmt>sla-profile>host-limits

Description This command limits the number of IPv4 DHCP hosts.

The **no** form of this command removes the number of IPv4 DHCP hosts from the SLA profile.

Parameters *max-nr-of-hosts* — Specifies the maximum number of IPv4 DHCP hosts.

Note: The operational maximum value may be smaller due to equipped hardware dependencies.

Values 0 to 131071

ipv4-overall

Syntax **ipv4-overall** *max-nr-of-hosts*
no ipv4-overall

Context config>subscr-mgmt>sla-profile>host-limits

Description This command limits the total number of IPv4 hosts.

The **no** form of this command removes the number of IPv4 hosts from the SLA profile.

Parameters *max-nr-of-hosts* — Specifies the maximum number of IPv4 hosts.

Note: The operational maximum value may be smaller due to equipped hardware dependencies.

Values 0 to 32767

ipv4-ppp

Syntax **ipv4-ppp** *max-nr-of-hosts*
no ipv4-ppp

Context config>subscr-mgmt>sla-profile>host-limits

Description This command limits the total number of IPv4 PPP hosts.

The **no** form of this command removes the number of IPv4 PPP hosts from the SLA profile.

Parameters *max-nr-of-hosts* — Specifies the maximum number of IPv4 PPP hosts.

Values 0 to 32767

ipv6-overall

Syntax **ipv6-overall** *max-nr-of-hosts*
no ipv6-overall

Context config>subscr-mgmt>sla-profile>host-limits

Description This command limits the total number of IPv6 hosts.

The **no** form of this command removes the number of IPv6 hosts from the SLA profile.

Parameters *max-nr-of-hosts* — Specifies the maximum number of IPv6 hosts.

Note: The operational maximum value may be smaller due to equipped hardware dependencies

Values 0 to 32767

ipv6-pd-ipoe-dhcp

Syntax **ipv6-pd-ipoe-dhcp** *max-nr-of-hosts*
no ipv6-pd-ipoe-dhcp

Context config>subscr-mgmt>sla-profile>host-limits

Description This command configures the total number of IPv6 DHCP PD hosts.

The **no** form of this command removes the number of IPv6 DHCP hosts from the SLA profile.

Parameters *max-nr-of-hosts* — Specifies the total number of IPv6 DHCP PD hosts.

Note: The operational maximum value may be smaller due to equipped hardware dependencies

Values 0 to 32767

ipv6-pd-overall

Syntax **ipv6-pd-overall** max-nr-of-hosts
no ipv6-pd-overall

Context config>subscr-mgmt>sla-profile>host-limits

Description This command limits the total number of IPv6-PD hosts.

The **no** form of this command removes the number of IPv6-PD hosts from the SLA profile.

Parameters *max-nr-of-hosts* — Specifies the maximum number of IPv6-PD hosts overall.

Note: The operational maximum value may be smaller due to equipped hardware dependencies

Values 0 to 32767

ipv6-pd-ppp-dhcp

Syntax **ipv6-pd-ppp-dhcp** *max-nr-of-hosts*
no ipv6-pd-ppp-dhcp

Context config>subscr-mgmt>sla-profile>host-limits

Description This command configures the maximum number of IPv6-WAN PPP DHCP hosts.

The **no** form of this command removes the number of IPv6-WAN PPP DHCP hosts from the SLA profile.

Parameters *max-nr-of-hosts* — Specifies the maximum number of IPv6-WAN PPP DHCP hosts.

Note: The operational maximum value may be smaller due to equipped hardware dependencies.

Values 0 to 32767

ipv6-wan-ipoe-dhcp

Syntax **ipv6-wan-ipoe-dhcp** *max-nr-of-hosts*
no ipv6-wan-ipoe-dhcp

Context config>subscr-mgmt>sla-profile>host-limits

Description This command configures the maximum number of IPv6-WAN PPP DHCP hosts.

The **no** form of this command removes the number of IPv6-WAN PPP DHCP hosts from the SLA profile.

Parameters *max-nr-of-hosts* — Specifies the maximum number of IPv6-WAN PPP DHCP hosts.

Note: The operational maximum value may be smaller due to equipped hardware dependencies

Values 0 to 32767

ipv6-wan-ipoe-slaac

Syntax **ipv6-wan-ipoe-slaac** *max-nr-of-hosts*
no ipv6-wan-ipoe-slaac

Context config>subscr-mgmt>sla-profile>host-limits

Description This command configures the maximum number of IPv6-WAN IPoE SLAAC hosts.

The **no** form of this command removes the number of IPv6-WAN IPoE SLAAC hosts from the SLA profile.

Parameters *max-nr-of-hosts* — Specifies the maximum number of IPv6-WAN IPoE SLAAC hosts.

Note: The operational maximum value may be smaller due to equipped hardware dependencies

Values 0 to 32767

ipv6-wan-overall

Syntax **ipv6-wan-overall** *max-nr-of-hosts*
no ipv6-wan-overall

Context config>subscr-mgmt>sla-profile>host-limits

Description This command configures the total number of IPv6 WAN hosts.

The **no** form of this command removes the number of IPV6 WAN hosts from the SLA profile.

Parameters *max-nr-of-hosts* — Specifies the maximum number of IPv6 WAN hosts.

Note: The operational maximum value may be smaller due to equipped hardware dependencies

Values 0 to 32767

ipv6-wan-ppp-dhcp

Syntax **ipv6-wan-ppp-dhcp** *max-nr-of-hosts*
no ipv6-wan-ppp-dhcp

Context config>subscr-mgmt>sla-profile>host-limits

Description This command configures the total number of IPv6 PPP DHCP WAN hosts.

The **no** form of this command removes the number of IPv6 PPP DHCP WAN hosts from the SLA profile.

Parameters *max-nr-of-hosts* — Specifies the maximum number of IPv6 PPP DHCP WAN hosts.

Note: The operational maximum value may be smaller due to equipped hardware dependencies

Values 0 to 32767

ipv6-wan-ppp-slaac

Syntax **ipv6-wan-ppp-slaac** *max-nr-of-hosts*
no ipv6-wan-ppp-slaac

Context config>subscr-mgmt>sla-profile>host-limits

Description This command configures the total number of SLAAC hosts.

The **no** form of this command removes the number of SLAAC hosts from the SLA profile.

Parameters *max-nr-of-hosts* — Specifies the maximum number of SLAAC hosts.

Note: The operational maximum value may be smaller due to equipped hardware dependencies

Values 0 to 32767

lac-overall

Syntax **lac-overall** *max-nr-of-hosts*
no lac-overall

Context config>subscr-mgmt>sla-profile>host-limits

Description This command configures the total number of L2TP LAC hosts

The **no** form of this command removes the number of L2TP LAC from the SLA profile.

Parameters *max-nr-of-hosts* — Specifies the maximum number of L2TP LAC hosts.

Note: The operational maximum value may be smaller due to equipped hardware dependencies

Values 0 to 32767

overall

Syntax **overall** *max-nr-of-hosts*
no overall

Context config>subscr-mgmt>sla-profile>host-limits

Description This command configures the total number of hosts.

The **no** form of this command reverts to the default.

Parameters *max-nr-of-hosts* — Specifies the maximum number of hosts.

Values 0 to 32767

remove-oldest

Syntax [**no**] **remove-oldest**

Context config>subscr-mgmt>sla-profile>host-limits

Description This command removes the oldest subscriber host when the host limit is reached.

The **no** form of this command maintains the oldest subscriber host when the host limit is reached.

bonding-selection

Syntax bonding-selection

Context config>subscr-mgmt>sla-profile>egress

Description This command enables the context to configure parameters belonging to this node for link selection behavior in a bonding context. These parameters are ignored outside of bonding subscribers.

hs-agg-rate-limit

Syntax **hs-agg-rate-limit** *kilobits-per-second*
no hs-agg-rate-limit

Context config>subscr-mgmt>sla-profile>egress

Description This command configures high scale (HS) aggregate rate limit of the SLA profile instance (SPI) associated with the subscriber in expanded SLA mode.

The aggregate rate of the subscriber (the primary shaper) is configured in the **config>subscr-mgmt>sub-profile>egress>hs-agg-rate-limit** context.

The **no** form of this command removes the value from the configuration.

Parameters *kilobits-per-second* — Specifies the HS egress aggregate rate limit.

Values 1 to 100000000

hs-low-burst-max-class

Syntax **hs-low-burst-max-class** *class*
no hs-low-burst-max-class

Context config>subscr-mgmt>sub-prof>egress

Description This command specifies which scheduling classes map to the low burst limit threshold of an egress HS primary shaper. The HS primary shaper is used to manage aggregate bandwidth of the subscriber with multiple SLA profile instances (expanded SLA mode).

Each HS primary shaper supports two burst thresholds, a low burst limit threshold and a high burst limit threshold. The two thresholds allow separation of burstiness between low and high scheduling classes.

When the low burst threshold of the HS primary shaper is reached, the lower scheduling classes, up to the scheduling class configured by this command, stops being served. Traffic on higher scheduling classes still goes through until the high burst threshold is reached. When the high burst threshold is exceeded, all scheduling classes associated with the HS primary shaper and stops being serve, which effectively shuts off the traffic flow.

Typically, the queues associated with higher scheduling classes are individually rate-limited so that their aggregate allowed throughput is less than the configured rate of the HS primary shaper. Determining the **hs-low-burst-max-class** *class* value involves anticipating the proper dividing line between the low and high scheduling classes by evaluating the forwarding behavior and SLA enforcement of each class.

By default, all scheduling classes are mapped to the low burst limit threshold. When mapping scheduling classes to the high burst limit threshold, an adequate value for the **config>card>fp>egress>hs-fixed-high-thresh-delta** should be specified (by default, it is set to 4000 bytes). This is because the queues associated with the lower scheduling classes may burst over the lower threshold during normal operation due to the scheduler forwarding whole packets. The **hs-fixed-high-thresh-delta** value should be set to at least two times the maximum frame size to prevent lower threshold class forwarding from also affecting the higher threshold classes when forwarding larger packet sizes. An insufficient high threshold delta defeats the intended purpose of mapping classes to the higher threshold.

The system uses the lowest value attainable for each low threshold aggregate burst limit without causing shaper underrun conditions. The high burst limit threshold is determined by adding the **hs-fixed-high-thresh-delta** value to the aggregate low burst limit threshold value.

The **hs-low-burst-max-class** value for HS primary shaper can be changed at any time in the subscriber profile (sub-profile).

The **no** form of the command restores the low burst limit threshold of the scheduling classes to the default value. This causes all scheduling classes associated with the HS primary shaper to be mapped to the low burst limit threshold.

Default 6

Parameters *class* — Specifies the highest scheduling class that is associated with the low burst limit threshold of the HS primary shaper. Scheduling classes that are higher than the scheduling class ID are associated with the high burst limit threshold.

Values 1 to 6

rate-thresholds

Syntax **rate-thresholds high** *high-percentage* **low** *low-percentage*
no rate-thresholds

Context config>subscr-mgmt>sla-profile>egress>bonding-selection

Description This command configures the rate thresholds that are required before decreasing or increasing the preferred link's weight with the specified change percentage.

The low threshold value must be lower than the high threshold value.

The **no** form of this command reverts to the default.

Default rate-thresholds high 90 low 80

Parameters *high-percentage* — Specifies the high threshold, as a percentage of the reference rate.

Values 56 to 99

low-percentage — Specifies the low threshold, as a percentage of the reference rate.

Values 55 to 98

weight

Syntax **weight** *initial-weight-percentage weight-change-percentage*
no weight

Context config>subscr-mgmt>sla-profile>egress>bonding-selection

Description This command configures the initial (and also the maximum) weight of the preferred connection and the value with which it can change.

The **no** form of this command resets the weight to the default.

Default weight 100 5

Parameters *initial-weight-percentage* — Specifies the initial or maximum weight, as a percentage of the offered weight.

Values 1 to 100

weight-change-percentage — Specifies the weight change, as a percentage.

Values 0 to 10

ip-filter

Syntax [**no**] **ip-filter** *filter-id*
no ip-filter [**force**]

Context config>subscr-mgmt>sla-profile>egress
config>subscr-mgmt>sla-profile>ingress

Description This command configures an egress or ingress IP filter.

The **no** form of this command reverts to the default.

Parameters *filter-id* — Specifies an existing IP filter policy ID.

Values 1 to 65535

SLA Profile QoS Commands

qos

Syntax **qos** *policy-id* [**vport-scheduler** | **port-scheduler**] [**force**]
no qos

Context config>subscr-mgmt>sla-prof>egress

Description This command specifies the egress QoS policy applicable to this SLA profile. The policy must already be defined in the **config>qos>sap-egress** context.

The **no** form of this command reverts to the default.

Default qos 1 port-scheduler

Parameters *policy-id* — Specifies the egress policy to be applied to the egress SLA profile.

Values 1 to 65535

vport-scheduler | port-scheduler — Specifies if a host queue with the port-parent option enabled should be scheduled within the context of a Vport port scheduler policy or a the port's port scheduler policy.

force — Forces a policy change

qos

Syntax **qos** *policy-id* [**shared-queuing** | **multipoint-shared** | **service-queuing**] [**force**]
no qos

Context config>subscr-mgmt>sla-prof>ingress

Description This command specifies the ingress QoS policy applicable to this SLA profile. The policy must already be defined in the **config>qos>sap-ingress** context.

The **no** form of this command reverts to the default.

Default qos 1

Parameters *policy-id* — Specifies the policy to be applied to the ingress SLA profile.

Values 1 to 65535

shared-queuing — Specifies the policy used by this SAP. When the value of this object is null it means that the SAP will use individual ingress QoS queues, instead of the shared ones.

multipoint-shared — This keyword is mutually exclusive with the **shared-queuing** and **service-queuing** keywords. When multipoint-shared is specified, the ingress forwarding plane will conserve hardware queues by performing two tier queuing on ingress unicast and multipoint packets through the SAP. Unicast service queues defined in the SAP ingress QoS policy are created for the SAP on the ingress forwarding plane without regard for the switch fabric destinations to which the SAP may need to forward (other destinations in the VPLS context). The multipoint queues defined in the SAP ingress QoS policy are not created for the SAP. Instead, all multipoint traffic is mapped to the unicast queues based on forwarding class in the first pass. In the second pass the unicast packets is mapped to the unicast shared queues while the multipoint traffic is mapped to the multipoint shared queues.

service-queuing — This keyword is mutually exclusive with the **multipoint-shared** and **shared-queuing** keywords to state that service queuing is needed.

force — Forces a policy change.

queue

Syntax [**no**] **queue** *queue-id*

Context config>subscr-mgmt>sla-prof>egress>qos
config>subscr-mgmt>sla-prof>ingress>qos

Description This command configures the context to configure egress or ingress queue parameters. Parameters defined in the **config>qos>sap-egress** *policy-id* or the **config>qos>sap-ingress** *policy-id* context are overridden by parameters specified in the subscriber management SLA profile context.

The classification and the queue mapping are shared by all the hosts on the same complex that use the same QoS policy (specified in the **sla-profile** SAP egress and SAP ingress policy IDs).

The queues are shared by all the hosts (of the same subscriber) on the same SAP that are using the same SLA profile. Queues are instantiated when, on a given SAP, a host of a subscriber is the first to use a certain SLA profile. This instantiation is referred to as an SLA profile instance.

The **no** form of this command removes the *queue-id* from the SLA profile.

Parameters *queue-id* — Specifies the queue ID for the SAP egress or ingress queue, expressed as a decimal integer. The ID uniquely identifies the queue within the profile.

Values 1 to 0

avg-frame-overhead

Syntax **avg-frame-overhead** *percent*
no avg-frame-overhead

Context config>subscr-mgmt>sla-prof>egress>qos>queue

Description This command configures the average frame overhead to define the average percentage that the offered load to a queue will expand during the frame encapsulation process before sending traffic on-the-wire. While the avg-frame-overhead value may be defined on any queue, it is only used by the system for queues that egress a SONET or SDH port or channel. Queues operating on egress Ethernet ports automatically calculate the frame encapsulation overhead based on a 20 byte per packet rule (8 bytes for preamble and 12 bytes for Inter-Frame Gap).

When calculating the frame encapsulation overhead for port scheduling purposes, the system determines the following values:

- Offered-load — The offered-load of a queue is calculated by starting with the queue depth in octets, adding the received octets at the queue and subtracting queue discard octets. The result is the number of octets the queue has available to transmit. This is the packet based offered-load.
- Frame encapsulation overhead — Using the avg-frame-overhead parameter, the frame encapsulation overhead is simply the queues current offered-load (how much has been received by the queue) multiplied by the avg-frame-overhead. If a queue had an offered load of 10000 octets and the avg-frame-overhead equals 10%, the frame encapsulation overhead would be 10000 x 0.1 or 1000 octets.

For egress Ethernet queues, the frame encapsulation overhead is calculated by multiplying the number of offered-packets for the queue by 20 bytes. If a queue was offered 50 packets then the frame encapsulation overhead would be 50 x 20 or 1000 octets.

- Frame based offered-load — The frame based offered-load is calculated by adding the offered-load to the frame encapsulation overhead. If the offered-load is 10000 octets and the encapsulation overhead was 1000 octets, the frame based offered-load would equal 11000 octets.
- Packet to frame factor — The packet to frame factor is calculated by dividing the frame encapsulation overhead by the queues offered-load (packet based). If the frame encapsulation overhead is 1000 octets and the offered-load is 10000 octets then the packet to frame factor would be 1000 / 10000 or 0.1. When in use, the avg-frame-overhead is the same as the packet to frame factor making this calculation unnecessary.
- Frame based CIR — The frame based CIR is calculated by multiplying the packet to frame factor with the queues configured CIR and then adding that result to that CIR. If the queue CIR is set at 500 octets and the packet to frame factor equals 0.1, the frame based CIR would be 500 x 1.1 or 550 octets.

- Frame based within-cir offered-load — The frame based within-cir offered-load is the portion of the frame based offered-load considered to be within the frame-based CIR. The frame based within-cir offered-load is the lesser of the frame based offered-load and the frame based CIR. If the frame based offered-load equaled 11000 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would be limited to 550 octets. If the frame based offered-load equaled 450 octets and the frame based CIR equaled 550 octets, the frame based within-cir offered-load would equal 450 octets (or the entire frame based offered-load).

As a special case, when a policer, queue or associated intermediate scheduler is configured with a CIR-weight equal to 0, the system automatically sets the queue's frame based within-cir offered-load to 0, preventing it from receiving bandwidth during the port scheduler's within-cir pass.

- Frame based PIR — The frame based PIR is calculated by multiplying the packet to frame factor with the queue's configured PIR and then adding the result to that PIR. If the queue PIR is set to 7500 octets and the packet to frame factor equals 0.1, the frame based PIR would be 7500 x 1.1 or 8250 octets.
- Frame based within-pir offered-load — The frame based within-pir offered-load is the portion of the frame based offered-load considered to be within the frame based PIR. The frame based within-pir offered-load is the lesser of the frame based offered-load and the frame based PIR. If the frame based offered-load equaled 11000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered-load would be limited to 8250 octets. If the frame based offered-load equaled 7000 octets and the frame based PIR equaled 8250 octets, the frame based within-pir offered load would equal 7000 octets.

Port scheduler operation using frame transformed rates — The port scheduler uses the frame based rates to determine the maximum rates that each queue may receive during the within-cir and above-cir bandwidth allocation passes. During the within-cir pass, a queue may receive up to its frame based within-cir offered-load. The maximum it may receive during the above-cir pass is the difference between the frame based within-pir offered load and the amount of actual bandwidth allocated during the within-cir pass.

SAP and subscriber SLA-profile average frame overhead override — The average frame overhead parameter on a sap-egress may be overridden at an individual egress queue basis. On each SAP and within the sla-profile policy used by subscribers an avg-frame-overhead command may be defined under the queue-override context for each queue. When overridden, the queue instance will use its local value for the average frame overhead instead of the sap-egress defined overhead.

The **no** form of this command reverts to the default. When set to 0, the system uses the packet based queue statistics for calculating port scheduler priority bandwidth allocation. If the no avg-frame-overhead command is executed in a queue-override queue id context, the avg-frame-overhead setting for the queue within the sap-egress QoS policy takes effect.

Default 0

Parameters *percent* — Specifies the average amount of packet-to-frame encapsulation overhead expected for the queue. This value is not used by the system for egress Ethernet queues.

Values 0 to 100, default

hs-queue-stat-mode

Syntax **hs-queue-stat-mode** *mode*
no hs-queue-stat-mode

Context config>subscr-mgmt>sla-prof>egress>qos

Description This command configures the mode of statistics collected for all the HS queues.

The **no** form of this command reverts to the default.

Default hs-queue-stat-mode no-override

Parameters *mode* — Specifies the egress HS queue stats mode.

Values no-override — Indicates no overrides are used.

v4-v6 — Indicates separate counters are collected for IPv4 and IPv6 instead of the normal queue statistics.

hs-wrr-group

Syntax [**no**] **hs-wrr-group** *group-id*

Context config>subscr-mgmt>sla-prof>egress>qos

Description This command configures the egress HS WRR group override parameters.

The **no** form of the command removes the group ID from the configuration.

Parameters *group-id* — Specifies the HS WRR group ID to override in the QoS policy table.

Values 1, 2

cbs

Syntax **cbs** *size-in-kbytes*
no cbs

Context config>subscr-mgmt>sla-prof>egress>qos>queue
config>subscr-mgmt>sla-prof>ingress>qos>queue

Description This command can be used to override specific attributes of the specified queue's CBS parameters. It is permissible, and possibly desirable, to oversubscribe the total CBS reserved buffers for a given access port egress buffer pool. Oversubscription may be desirable due to the potential large number of service queues and the economy of statistical multiplexing the individual queues' CBS settings into the defined reserved total.

When oversubscribing the reserved total, it is possible for a queue depth to be lower than its CBS setting and still not receive a buffer from the buffer pool for an ingress frame. As more queues are using their CBS buffers and the total in use exceeds the defined reserved total, essentially the buffers are being removed from the shared portion of the pool without the shared in use average and total counts being decremented. This can affect the operation of the high and low priority RED slopes on the pool, causing them to miscalculate when to start randomly drop packets.

The **no** form of this command returns the CBS size to the size as configured in the QoS policy.

Default no cbs

Parameters *size-in-kbytes* — The size parameter is an integer expression of the number of kilobytes reserved for the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that no reserved buffers are required by the queue (a minimal reserved size can still be applied for scheduling purposes).

Values 0 to 1048576, **default**

high-prio-only

Syntax **high-prio-only** *percent*
no high-prio-only

Context config>subscr-mgmt>sla-prof>egress>qos>queue
config>subscr-mgmt>sla-prof>ingress>qos>queue

Description This command configures the value of the percentage of buffer space for the queue, used exclusively by high priority packets. The specified value overrides the default value for the context.

The priority of a packet can only be set in the SAP ingress QoS policy and is only applicable on the ingress queues for a SAP. The **high-prio-only** parameter is used to override the default value derived from the **network-queue** command.

The defined **high-prio-only** value cannot be greater than the MBS size of the queue. Attempting to change the MBS to a value smaller than the high priority reserve will generate an error and fail execution. Attempting to set the **high-prio-only** value larger than the current MBS size will also result in an error and fail execution.

The **no** form of this command returns high-prio-only to the size as configured in the QoS policy.

Parameters *percent* — The *percent* parameter is the percentage reserved for high priority traffic on the queue. If a value of 10KBytes is desired, enter the value 10. A value of 0 specifies that none of the MBS of the queue is reserved for high priority traffic. This does not affect RED slope operation for packets attempting to be queued.

Values 0 to 100, default

hs-class-weight

Syntax **hs-class-weight** *weight*
no hs-class-weight

Context config>subscr-mgmt>sla-prof>egress>qos>queue
config>subscr-mgmt>sla-prof>egress>qos>hs-wrr-grp

Description This command configures the class-weight override for expanded egress HS queues or the WRR group.

The **no** form of this command removes the weight value from the configuration.

Parameters *weight* — Specifies the weight of the scheduling class.

Values 1, 2, 4, 8

rate

Syntax **rate** *rate*
no rate

Context config>subscr-mgmt>sla-prof>egress>qos>hs-wrr-grp

Description This command configures the rate (PIR) override for the WRR group.

The **no** form of this command removes the rate from the configuration.

Parameters *rate* — Specifies the PIR expressed as a percentage of line rate in Kb/s.

Values 1 to 2000000000, max

hs-wred-queue-policy

Syntax **hs-wred-queue-policy** *name*
no hs-wred-queue-policy

Context config>subscr-mgmt>sla-prof>egress>qos>queue

Description This command specifies the name of the slope-policy override to be applied for the HS queue of this SLA profile instance.

The **no** form of this command removes the policy name from the configuration.

Parameters *name* — Specifies the policy name up to 32 characters.

hs-wrr-weight

Syntax **hs-wrr-weight** *weight*
no hs-wrr-weight

Context config>subscr-mgmt>sla-prof>egress>qos>queue

Description This command configures the SLA profile instance WRR weight override for the HS queue. When a weight value is not specified, there is no override, meaning, the WRR weight is taken from the sap-egress policy.

The **no** form of this command removes the weight value from the configuration.

Parameters *weight* — Specifies the class-weight override for expanded egress HS queues.

Values 1 to 127

mbs

Syntax **mbs** *size* [**bytes** | **kilobytes**]
no mbs

Context config>subscr-mgmt>sla-prof>egress>qos>queue

Description This command configures the maximum size for the queue.

The sum of the MBS for all queues on an egress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer is available when needed or that the packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

The **no** form of this command returns the MBS size to the size as configured in the QoS policy.

Parameters *size* — This required parameter specifies that the MBS is expressed as an integer representing the required size in either bytes or kilobytes. The default is **kilobytes**. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly to define the size. By specifying the keyword **default** sets the MBS to its default value.

Values 0 to 1073741824, **default**

bytes — Specifies that the value given for *size* is interpreted as the queue's MBS value in bytes.

kilobytes — Specifies that the value given for *size* is interpreted as the queue's MBS value in kilobytes.

mbs

Syntax **mbs** *size* [**bytes** | **kilobytes**]
no mbs

Context config>subscr-mgmt>sla-prof>ingress>qos>queue

Description The Maximum Burst Size (MBS) command configures the explicit definition of the maximum amount of buffers allowed for a specific queue.

The MBS value is used by a queue to determine whether it has exhausted all of its buffers while enqueuing packets. Once the queue has exceeded the amount of buffers allowed by MBS, all packets are discarded until packets have been drained from the queue.

The sap-ingress context for mbs provides a mechanism for overriding the default maximum size for the queue.

The sum of the MBS for all queues on an ingress access port can oversubscribe the total amount of buffering available. When congestion occurs and buffers become scarce, access to buffers is controlled by the RED slope a packet is associated with. A queue that has not exceeded its MBS size is not guaranteed that a buffer is available when needed or that the packet's RED slope will not force the discard of the packet. Setting proper CBS parameters and controlling CBS oversubscription is one major safeguard to queue starvation (when a queue does not receive its fair share of buffers). Another is properly setting the RED slope parameters for the needs of services on this port or channel.

The **no** form of this command returns the MBS size to the size as configured in the QoS policy.

Parameters *size* — This required parameter specifies that the MBS is expressed as an integer representing the required size in either bytes or kilobytes. The default is **kilobytes**. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly to define the size. By specifying the keyword **default** sets the MBS to its default value.

Values 0 to 1073741824, **default**

bytes — Specifies that the value given for *size* is interpreted as the queue's MBS value is in bytes.

kilobytes — Specifies that the value given for *size* is interpreted as the queue's MBS value is in kilobytes.

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

Context config>subscr-mgmt>sla-prof>egress>qos>queue

Description This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile then out-of-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent command's *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of this command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Parameters *pir-rate* — Defines the administrative PIR rate, in kb/s, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.
Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queues **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 to 2000000000, max

Default max

cir-rate — Defines the administrative CIR rate, kb/s, for the queue. The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.
Fractional values are not allowed and must be given as a positive integer.

Values 0 to 2000000000, max

Default 0

qos-marking-from-sap

Syntax [**no**] **qos-marking-from-sap**

Context config>subscr-mgmt>sla-profile>egress

Description This command sets the QoS policy from which the egress QoS marking rules are applied.

Note: If applied to a managed SAP, the default s-egress qos-policy (sap-egress 1) cannot be changed.

The **no** form of this command reverts to the egress QoS marking defined in SAP-egress policy defined at sla-profile level.

Default qos-marking-from-sap

report-rate

Syntax **report-rate agg-rate-limit**
report-rate scheduler *scheduler-name*
report-rate pppoe-actual-rate
report-rate policer *policer-id*
report-rate rfc5515-actual-rate
no report-rate

Context config>subscr-mgmt>sla-prof>ingress
config>subscr-mgmt>sla-prof>egress

Description This command configures the source for Tx and Rx connect speeds in AVP 38 (Rx Connect Speed) and AVP 24 (Tx Connect Speed) of an L2TP session established on a LAC.

The **no** form of this command reverts to the default.

Parameters **agg-rate-limit** — Specifies that the rate (egress only) is taken from:

1. the agg-rate RADIUS override (RADIUS VSA “Alc-Subscriber-QoS-Override” in a RADIUS Access-Accept message) if present
2. the configured agg-rate-limit in the **config>subscr-mgmt>sub-prof>egr** context
3. fall back to the default (no report-rate)

scheduler-name — Specifies the rate taken from the **scheduler** *scheduler-name* up to 32 characters. If the **scheduler** *scheduler-name* is not present in the scheduler-policy configured in the **config>subscr-mgmt>sub-prof>egr** context, fall back to the default (no report-rate).

pppoe-actual-rate — Specifies rates taken from the DSL Line characteristics PPPoE tags (Actual Data Rate Upstream/Downstream) if present; otherwise fall back to the default (no report-rate).

policer-id — Specifies the rate taken from the policer with the specified ID.

rfc5515-actual-rate — Puts the same value as the transmitted Actual-Data-Rate-Upstream AVP in the Rx-Connect-Speed AVP, and the same value as the transmitted Actual-Data-Rate-Downstream AVP in the Tx-Connect-Speed AVP

scheduler-policy

Syntax **scheduler-policy** *scheduler-name* **rate** *pir-rate* [**cir** *cir-rate*]
no scheduler-policy

Context config>subscr-mgmt>sla-prof>egress

Description This command specifies a scheduler policy to associate to the sla profile. Scheduler policies are configured in the **configure>qos>scheduler>policy** context. Each scheduler policy is divided up into groups of schedulers based on the tier each scheduler is created under. A tier is used to give structure to the schedulers within a policy and define rules for parent scheduler associations. The policy defines the hierarchy and operating parameters for virtual schedulers.

The **no** form of this command removes the scheduler-policy-name from the configuration.

Parameters *scheduler-policy-name* — Specifies an existing scheduler policy name up to 32 characters.

scheduler-policy

Syntax **scheduler-policy** *scheduler-name* **rate** *pir-rate* [**cir** *cir-rate*]
no scheduler-policy *scheduler-name*

Context config>subscr-mgmt>sla-prof>egress>scheduler-policy

Description This command provides a way to override parameters of the existing scheduler associated with the egress scheduler policy. A scheduler defines bandwidth controls that limit each child (other schedulers and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have queues or other schedulers defined as child associations. The scheduler can be a child (take bandwidth from a scheduler in a higher tier).

The **no** form of this command reverts to the default.

Parameters *scheduler-name* — Specify an existing scheduler policy name up to 32 characters.

pir-rate — Specifies the PIR rate in kb/s. This parameter overrides the administrative PIR used by the scheduler. When the rate command is executed, a valid PIR setting must be explicitly defined.
Fractional values are not allowed and must be given as a positive integer.

Values 1 to 3200000000, max

cir-rate — Specifies the CIR rate in kb/s. This parameter overrides the administrative CIR used by the scheduler. When the rate command is executed, a CIR setting is optional. The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers or queues.
Fractional values are not allowed and must be given as a positive integer.

Values 0 to 3200000000, sum, max

Default sum

use-ingress-l2tp-dscp

Syntax [**no**] **use-ingress-l2tp-dscp**

Context config>subscr-mgmt>sla-prof>egress

Description This command enables the use of the DSCP marking taken from the L2TP header received on an L2TP Access Concentrator (LAC) for egress classification for the subscriber host using the associated sla-profile.

This command is ignored if the ingress packet is not identified as an L2TP packet.

The **no** form of this command reverts to the default.

one-time-http-redirection

Syntax **one-time-http-redirection ip-filter** *filter-id*
one-time-http-redirection

Context config>subscr-mgmt>sla-prof

Description This command specify the one-time http redirection filter id. This filter will apply to the host when host is created, and is replaced by the sla-profile ingress filter (configured in the **config>subscr-mgmt>sla-prof>ingress** context) after first HTTP request from host has been redirected.

Note: The system does not check if the configured filter include http-redirection entry. If the filter does not include the http-redirection then it will not be replaced in future.

If 7750 SR receives filter insertion via CoA or access-accept when one-time redirection filter is still active then the received filter entries will only be applied to the sla-profile ingress filter. And after 1st http redirection, the original sla-profile ingress filter + received filter will replace the redirection filter.

The **no** form of this command reverts to the default.

Parameters *filter-id* — Specifies the ID of filter that is used for HTTP redirection.

Values 1 to 65535

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

Context config>subscr-mgmt>sla-prof>ingress>qos>queue

Description This command defines the administrative Peak Information Rate (PIR) and the administrative Committed Information Rate (CIR) parameters for the queue. The PIR defines the maximum rate that the queue can transmit packets through the switch fabric (for SAP ingress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. For SAP ingress, the CIR also defines the rate that packets are considered in-profile by the system. In-profile then out-of-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent command's *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at anytime, altering the PIR and CIR rates for all queues created through the association of the SAP ingress or SAP egress QoS policy with the *queue-id*.

The **no** form of this command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Parameters *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.

Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queues **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 to 2000000000, max

Default max

cir-rate — Specifies the **cir** parameter used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.

Fractional values are not allowed and must be given as a positive integer.

Values 0 to 2000000000, max

Default 0

policer

Syntax [**no**] **policer** *policer-id*

Context config>subscr-mgmt>sla-prof>ingress>qos
config>subscr-mgmt>sla-prof>egress>qos
config>subscr-mgmt>sub-profile>hsmda>ingress-qos>qos>policer

Description This command is used in the sap-ingress and sap-egress QoS policies to create, modify or delete a policer. Policers are created and used in a similar manner to queues. The policer ID space is separate from the queue ID space, allowing both a queue and a policer to share the same ID. The sap-ingress policy may have up to 32 policers (numbered 1 through 32) may be defined while the sap-egress QoS policy supports a maximum of 8 (numbered 1 through 8). While a policer may be defined within a QoS policy, it is not actually created on SAPs or subscribers associated with the policy until a forwarding class is mapped to the policer's ID.

All policers must be created within the QoS policies. A default policer is not created when a sap-ingress or sap-egress QoS policy is created.

Once a policer is created, the policer's metering rate and profiling rates may be defined as well as the policer's maximum and committed burst sizes (MBS and CBS respectively). Unlike queues which have dedicated counters, policers allow various stat-mode settings that define the counters that is associated with the policer. Another supported feature—packet-byte-offset—provides a policer with the ability to modify the size of each packet based on a defined number of bytes.

Once a policer is created, it cannot be deleted from the QoS policy unless any forwarding classes that are mapped to the policer are first moved to other policers or queues.

The system will allow a policer to be created on a SAP QoS policy regardless of the ability to support policers on objects where the policy is currently applied. The system only scans the current objects for policer support and sufficient resources to create the policer when a forwarding class is first mapped to the policer ID. If the policer cannot be created due to one or more instances of the policy not supporting policing or having insufficient resources to create the policer, the forwarding class mapping will fail.

The **no** form of this command is used to delete a policer from a sap-ingress or sap-egress QoS policy. The specified policer cannot currently have any forwarding class mappings for the removal of the policer to succeed. It is not necessary to actually delete the policer ID for the policer instances to be removed from SAPs or subscribers associated with the QoS policy once all forwarding classes have been moved away from the policer. It is automatically deleted from each policing instance although it still appears in the QoS policy.

Parameters *policer-id* — Specifies the policer ID. The policer-id must be specified when executing the policer command. If the specified ID already exists, the system enters that policer's context to allow the policer's parameters to be modified. If the ID does not exist and is within the allowed range for the QoS policy type, a context for the policer ID is created (depending on the system's current create keyword requirements which may require the create keyword to actually add the new policer ID to the QoS policy) and the system will enter that new policer's context for possible parameter modification.

Values 1 to 63

cbs

Syntax **cbs** *size* [**bytes** | **kilobytes**]
no cbs

Context config>subscr-mgmt>sla-prof>ingress>qos>policer
config>subscr-mgmt>sla-prof>egress>qos>policer

Description This command is used to configure the policer's CIR leaky bucket's exceed threshold. The CIR bucket's exceed threshold represents the committed burst tolerance allowed by the policer. If the policer's forwarding rate is equal to or less than the policer's defined CIR, the CIR bucket depth hovers around the 0 depth with spikes up to the maximum packet size in the offered load. If the forwarding rate increases beyond the profiling rate, the amount of data allowed to be in-profile above the rate is capped by the threshold.

The policer's **cbs** size defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** form of this command returns the policer to its default CBS size.

Parameters *size* — Specifies the size parameter and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

bytes — Specifies the size parameter the size parameter in bytes. When **bytes** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

kilobytes — Specifies the size parameter in kilobytes. When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

Values 0 to 16777216

Default kilobyte

mbs

Syntax **mbs** *size* [**bytes** | **kilobytes**]
no mbs

Context config>subscr-mgmt>sla-prof>ingress>qos>policer
config>subscr-mgmt>sla-prof>egress>qos>policer

Description This command configures the MBS for the QoS policer.

Parameters *size* — This required parameter specifies that the MBS is expressed as an integer representing the required size in either bytes or kilobytes. The default is **kilobytes**. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly to define the size. By specifying the keyword **default** sets the MBS to its default value.

Values 0 to 16777216, **default**

bytes — Specifies that the value given for *size* is interpreted as the queue's MBS value in bytes.

kilobytes — Specifies that the value given for *size* is interpreted as the queue's MBS value in kilobytes.

mbs

Syntax **mbs** *size* [**bytes** | **kilobytes**]
no mbs

Context config>subscr-mgmt>sub-profile>hsmda>egress-qos>qos>queue

Description This command overrides definition of the maximum amount of buffers allowed for a specific queue. The value is given in bytes or kilobytes and overrides the default value for the context.

The **no** form of this command returns the queue to its default MBS size.

Parameters *size* [**bytes** | **kilobytes**] — Specifies the and is expressed as an integer representing the required size in either bytes or kilobytes. The default is kilobytes. The optional **byte** and **kilobyte** keywords are mutually exclusive and are used to explicitly define whether size represents bytes or kilobytes.

byte — When **byte** is defined, the value given for size is interpreted as the queue's MBS value given in bytes.

Values 0 to 2688000

kilobyte — When **kilobytes** is defined, the value is interpreted as the queue's MBS value given in kilobytes.

Values 0 to 2625

Default kilobyte

packet-byte-offset

Syntax **packet-byte-offset** {**add** *bytes* | **subtract** *bytes*}
no packet-byte-offset

Context config>subscr-mgmt>sla-prof>ingress>qos>policer
config>subscr-mgmt>sla-prof>egress>qos>policer

Description This command is used to modify the size of each packet handled by the policer by adding or subtracting a number of bytes. The actual packet size is not modified; only the size used to determine the bucket depth impact is changed. The **packet-byte-offset** command is meant to be an arbitrary mechanism the can be used to either add downstream frame encapsulation or remove portions of packet headers. Both the policing metering and profiling throughput is affected by the offset as well as the stats associated with the policer.

When child policers are adding to or subtracting from the size of each packet, the parent policer's **min-thresh-separation** value should also need to be modified by the same amount.

The policer's **packet-byte-offset** defined in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied. Packet byte offset settings are not included in the applied rate when (queue) frame based accounting is configured and the policer is managed by HQoS, however the offsets are applied to the statistics.

The **no** form of this command removes the per packet size modifications from the policer.

Parameters **add** *bytes* — Specifies the packet byte offset. The **add** keyword is mutually exclusive to the **subtract** keyword. Either **add** or **subtract** must be specified. When **add** is defined the corresponding bytes parameter specifies the number of bytes that is added to the size each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is increased by the amount being added to the size of each packet.

Values 0 to 31

subtract *bytes* — Specifies the packet byte offset. The **subtract** keyword is mutually exclusive to the **add** keyword. Either **add** or **subtract** must be specified. When b is defined the corresponding bytes parameter specifies the number of bytes that is subtracted from the size of each packet associated with the policer for rate metering, profiling and accounting purposes. From the policer's perspective, the maximum packet size is reduced by the amount being subtracted from the size of each packet.

Note: The minimum resulting packet size used by the system is 1 byte.

Values ingress 1 to 32
egress: 1 to 64

rate

Syntax **rate** {**max** | *rate*} [**cir** {**max** | *rate*}]

Context config>subscr-mgmt>sla-prof>ingress>qos>policer
config>subscr-mgmt>sla-prof>egress>qos>policer

Description This command is used to configure the policer's metering and optional profiling rates. The metering rate is used by the system to configure the policer's PIR leaky bucket's decrement rate while the profiling rate configures the policer's CIR leaky bucket's decrement rate. The decrement function empties the bucket while packets applied to the bucket attempt to fill it based on the each packets size. If the bucket fills faster than how much is decremented per packet, the bucket's depth eventually reaches it's exceed (CIR) or violate (PIR) threshold. The **cbs**, **mbs**, and **high-prio-only** commands are used to configure the policer's PIR and CIR thresholds.

If a packet arrives at the policer while the bucket's depth is less than the threshold associated with the packet, the packet is considered to be conforming to the bucket's rate. If the bucket depth is equal to or greater than the threshold, the packet is considered to be in the exception state. For the CIR bucket, the exception state is exceeding the CIR rate while the PIR bucket's exception state is violating the PIR bucket rate. If the packet is violating the PIR, the packet is marked red and is discarded. If the packet is not red, it may be green or yellow based on the conforming or exceeding state from the CIR bucket.

When a packet is red neither the PIR or CIR bucket depths are incremented by the packets size. When the packet is yellow the PIR bucket is incremented by the packet size, but the CIR bucket is not. When the packet is green, both the PIR and CIR buckets are incremented by the packet size. This ensures that conforming packets impact the bucket depth while exceeding or violating packets do not.

The policer's **adaptation-rule** command settings are used by the system to convert the specified rates into hardware timers and decrement values for the policer's buckets.

By default, the policer's metering rate is **max** and the profiling rate is 0 kb/s (all packets out-of-profile).

The **rate** settings defined for the policer in the QoS policy may be overridden on an **sla-profile** or SAP where the policy is applied.

The **no** form of this command restores the default metering and profiling rate to a policer.

Parameters {**max** | *rate*} — Specifies the packet byte offset. Specifying the keyword **max** or an explicit kilobits-per-second parameter directly following the rate command is required and identifies the policer's metering rate for the PIR leaky bucket. When the policer is first created, the metering rate defaults to max. The kilobits-per-second value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second. When max is specified, the maximum policer rate used is equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the PIR used is equivalent to max.

Values max or 1 to 2000000000

cir {**max** | *rate*} — Specifies the packet byte offset. The optional **cir** keyword is used to override the default CIR rate of the policer. Specifying the keyword max or an explicit kilobits-per-second parameter directly following the cir keyword is required and identifies the policer's profiling rate for the CIR leaky bucket. When the policer is first created, the profiling rate defaults to 0 kb/s. The kilobits-per-second value must be expressed as an integer and defines the rate in kilobits-per-second. The integer value is multiplied by 1,000 to derive the actual rate in bits-per-second. When max is specified, the maximum policer rate used is equal to the maximum capacity of the card on which the policer is configured. If the policer rate is set to a value larger than the maximum rate possible for the card, then the CIR used is equivalent to max.

Values max or 0 to 2000000000

stat-mode

Syntax **stat-mode** *stat-mode*
no stat mode

Context config>subscr-mgmt>sla-prof>ingress>qos>policer
config>subscr-mgmt>sla-prof>ingress>qos>queue
config>subscr-mgmt>sla-prof>egress>qos>policer
config>subscr-mgmt>sla-prof>egress>qos>queue

Description This command is used to configure the forwarding plane octet and packet counters of a policer or queue to count packets of a specific type or state. For example separate counters for IPv4/IPv6 or separate counters for offered high and low priority policed traffic.

For policers, this command overrides the policer stat-mode configuration as defined in the sap-ingress or sap-egress qos policy. For details on sap-ingress and sap-egress policer stat-mode, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Quality of Service Guide*. For use in Enhanced Subscriber Management (ESM) context only, an additional stat-mode enables separate counters for IPv4 and IPv6 packets.

When a policer's stat-mode is changed while the sla profile is in use, any previous counter values are lost and any new counters are set to zero.

For queues, this command sets the stat-mode. Queue stat-mode is only available for use in Enhanced Subscriber Management (ESM) context to enable separate IPv4/IPv6 counters.

A queue's stat-mode cannot be changed while the SLA profile is in use.

The **no** form of this command reverts to the default.

Default For policers, the default is no stat-mode override. The sap-ingress or sap-egress stat-mode is used instead.

For queues, the default is to count in-/out-of-profile octets and packets.

Parameters *stat-mode* — Specifies the stat mode for the policer.

For ingress and egress qos queue stat-mode overrides.

For ingress and egress qos policer stat-mode overrides, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Quality of Service Guide* for details on the **sap-ingress** and **sap-egress policer stat-mode** parameters.

For use in Enhanced Subscriber Management (ESM) context only:

Values no-stats, minimal, offered-profile-no-cir, offered-priority-no-cir, offered-profile-cir, offered-priority-cir, offered-total-cir, offered-limited-profile-cir, offered-profile-capped-cir, offered-limited-capped-cir, v4-v6 (count IPv4 and IPv6 forwarded/dropped octets and packets separately)

9.25.2.1.23 Subscriber Identification Policy Commands

sub-ident-policy

Syntax [**no**] **sub-ident-policy** *sub-ident-policy-name*

Context config>subscr-mgmt

Description This command configures a subscriber identification policy. Each subscriber identification policy can have a default subscriber profile defined. The subscriber identification policy default subscriber profile overrides the system default and the subscriber SAP default subscriber profiles. Defining a subscriber identification policy default subscriber profile is optional.

The subscriber identification policy default subscriber profile cannot be defined with the subscriber profile name default.

Defining a subscriber profile as a subscriber identification policy default subscriber profile will cause all active subscribers currently associated with a subscriber SAP using the policy and associated with a subscriber policy through the system default or subscriber SAP default subscriber profiles to be reassigned to the subscriber policy defined as default on the subscriber identification policy.

Attempting to delete a subscriber profile that is currently defined as a default for a subscriber identification policy will fail.

When attempting to remove a subscriber identification policy default subscriber profile definition, the system will evaluate each active subscriber on all subscriber SAPs the subscriber identification policy is currently associated with that are using the default definition to determine whether the active subscriber can be either reassigned to a subscriber SAP default or the system default subscriber profile. If all active subscribers cannot be reassigned, the removal attempt will fail.

The **no** form of this command reverts to the default.

Parameters *sub-ident-policy-name* — Specifies the name of the subscriber identification policy, up to 32 characters.

sub-ident-policy

Syntax **sub-ident-policy** *sub-ident-policy-name*

Context config>service>vpls>sap>sub-sla-mgmt
config>service>ies>if>sap>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt

Description This command associates a subscriber identification policy to this SAP. The subscriber identification policy must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sub-ident-policy** context.

Subscribers are managed by the system through the use of subscriber identification strings. A subscriber identification string uniquely identifies a subscriber. For static hosts, the subscriber identification string is explicitly defined with each static subscriber host.

For dynamic hosts, the subscriber identification string must be derived from the DHCP ACK message sent to the subscriber host. The default value for the string is the content of Option 82 CIRCUIT-ID and REMOTE-ID fields interpreted as an octet string. As an option, the DHCP ACK message may be processed by a subscriber identification policy which has the capability to parse the message into an alternative ASCII or octet string value.

When multiple hosts on the same port are associated with the same subscriber identification string they are considered to be host members of the same subscriber.

The **no** form of this command removes the default subscriber identification policy from the SAP configuration.

Parameters *sub-ident-policy-name* — Specifies a subscriber identification policy for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sub-ident-policy** context.

app-profile-map

Syntax **app-profile-map**

Context config>subscr-mgmt>sub-ident-pol

Description This command enables the context to configure an application profile mapping.

entry

Syntax **entry key** *app-profile-string* **app-profile** *app-profile-name*
no entry key *app-profile-string*

Context config>subscr-mgmt>sub-ident-pol>app-profile-map

Description This command configures an application profile string.

The **no** form of this command removes the values from the configuration.

Parameters *app-profile-string* — Specifies the application profile string up to 16 characters.

app-profile-name — Specifies the application profile name up to 32 characters.

use-direct-map-as-default

Syntax [**no**] **use-direct-map-as-default**

Context config>subscr-mgmt>sub-ident-pol>app-profile-map
config>subscr-mgmt>sub-ident-pol>sla-profile-map

Description This command enables direct mapping of application profile as default. With this flag, a script returned string is used as the named profile. If the named profiled cannot be found, the default profile is used.

The **no** form of this command disables the direct mapping.

primary

Syntax **primary**

Context config>subscr-mgmt>sub-ident-pol

Description This command enables the context to configure primary identification script parameters.

secondary

Syntax **secondary**

Context config>subscr-mgmt>sub-ident-pol

Description This command enables the context to configure secondary identification script parameters.

tertiary

Syntax **tertiary**

Context config>subscr-mgmt>sub-ident-pol

Description This command enables the context to configure tertiary identification script parameters.

script-url

Syntax **script-url** *dhcp-script-url*
no script-url

Context config>subscr-mgmt>sub-ident-pol>primary
config>subscr-mgmt>sub-ident-pol>secondary
config>subscr-mgmt>sub-ident-pol>tertiary

Description This command specifies the URL of the identification scripts.

The **no** form of this command reverts to the default.

Parameters *dhcp-primary-script-url* — Specifies the URL of the primary identification script up to 180 characters.

dhcp-secondary-script-url — Specifies the URL of the secondary identification script up to 180 characters.

dhcp-tertiary-script-url — Specifies the URL of the tertiary identification script up to 180 characters.

sla-profile-map

Syntax **sla-profile-map**

Context config>subscr-mgmt>sub-ident-pol

Description This command enables the context to configure SLA profile mapping parameters.

sub-profile-map

Syntax **sla-profile-map**

Context config>subscr-mgmt>sub-ident-pol

Description This command enables the context to configure subscriber profile mapping parameters.

entry

Syntax **entry key** *sla-profile-string* **sla-profile** *sla-profile-name*
no entry key *sla-profile-string*

Context config>subscr-mgmt>sub-ident-pol>sla-profile-map
config>subscr-mgmt>sub-prof>sla-profile-map

Description This command configures an SLA profile string. Each subscriber identification string can be provisioned into a subscriber mapping table providing an explicit mapping of the string to a specific subscriber profile. This allows certain subscribers to be directly mapped to the appropriate subscriber profile in the event that the default mappings are not desired for the subscriber.

An explicit mapping of a subscriber identification string to a subscriber profile cannot be defined with the subscriber profile name default. It is possible for the subscriber identification string to be entered in the mapping table without a defined subscriber profile which can result in the explicitly defined subscriber to be associated with the subscriber profile named default.

Explicitly mapping a subscriber identification string to a subscriber profile will cause an existing active subscriber associated with the string to be reassigned to the newly mapped subscriber profile. An explicit mapping overrides all default subscriber profile definitions.

Attempting to delete a subscriber profile that is currently defined as in an explicit subscriber identification string mapping will fail.

The system will fail the removal attempt of an explicit subscriber identification string mapping to a subscriber profile definition when an active subscriber is using the mapping and cannot be reassigned to a defined default non-provisioned subscriber profile.

The **no** form of this command reverts to the default.

Parameters *sla-profile-string* — Identifies the SLA profile string, up to 16 characters.

sla-profile-name — Identifies the SLA profile name, up to 32 characters.

entry

Syntax **entry key** *sub-profile-string* **sub-profile** *sub-profile-name*
no entry key *sub-profile-string*

Context config>subscr-mgmt>sub-ident-pol>sub-profile-map

Description This command configures a subscriber profile string.

The **no** form of this command reverts to the default.

Parameters *sub-profile-string* — Specifies the subscriber profile string, up to 16 characters.

sub-profile-name — Specifies the subscriber profile name, up to 32 characters.

9.25.2.1.24 Subscriber Profile Commands

accu-stats-policy

Syntax **accu-stats-policy** *policy-name* [**create**]
no accu-stats-policy *policy-name*

Context config>subscr-mgmt

Description This command creates a storage policy for cumulative statistics for subscribers. The policy defines the specific direction for the policer or the queue to be stored and performs the following functions.

- The policy stores subscriber statistics even if the subscriber session has ended. The subscriber statistics can be viewed even if the subscriber is offline.
- When the subscriber session ends, the statistics are added to the past statistics stored in memory so that all previous session statistics are accumulated. The accumulated statistics are reset to zero when the chassis reboots.

The **no** form of this command deletes the policy only when it is no longer referenced by a subscriber profile.

Parameters *policy-name* — Specifies the name for the policy, up to 32 characters.

create — Configures an entry for the policy.

entry

Syntax [**no**] **entry direction** *direction* **type** *type* **id** *id*

Context config>subscr-mgmt>accu-stats-policy

Description This command defines the direction of the policer or queue to the stored and accumulated policy.

The **no** form of this command removes the entry.

Parameters *direction* — Specifies the direction of the queue or policer.

Values egress, ingress

type — Specifies whether the entry is for a queue or policer.

Values queue, policer

id — Specifies the queue or policer ID.

Values 1 to 63

sub-profile

Syntax [**no**] **sub-profile** *subscriber-profile-name*

Context config>subscr-mgmt

Description This command enables the context to configure a subscriber profile. A subscriber profile is a template used to define the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for subscribers using the subscriber profile. Subscriber profiles also allow for specific SLA profile definitions when the default definitions from the subscriber identification policy must be overridden.

Subscribers are either explicitly mapped to a subscriber profile template or are dynamically associated by one of various non-provisioned subscriber profile definitions.

A subscriber host can be associated with a subscriber profile in the following ways, listed from lowest to highest precedence:

1. The subscriber profile named default.
2. The subscriber profile defined as the subscriber SAP default.
3. The subscriber profile found by the subscriber identification policy sub-profile-map.
4. The subscriber profile found by the subscriber identification policy explicit map.

In the event that no defaults are defined and the subscriber identification string is not explicitly provisioned to map to a subscriber profile, either the static subscriber host creation will fail or the dynamic subscriber host DHCP ACK is discarded.

Default Subscriber profile:

When a subscriber profile is created with the *subscriber-profile-name* default, it is used when no other subscriber profile is associated with the subscriber host by the system. Creating a subscriber profile with the *subscriber-profile-name* default is optional. If a default subscriber profile is not created, all subscriber hosts subscriber identification strings must match either a non-provisioned default or be provisioned as an explicit match to a subscriber profile.

The default profile has no effect on existing active subscriber on the system as they exist due to higher precedence mappings.

Attempting to delete any subscriber profile (including the profile named default) while in use by existing active subscribers will fail.

The **no** form of this command reverts to the default.

Parameters *subscriber-profile-name* — Specifies the name of the subscriber profile, up to 32 characters.

accounting-policy

Syntax **accounting-policy** *acct-policy-id*
no accounting-policy

Context config>subscr-mgmt>sub-prof

Description This command specifies the policy to use to collect accounting statistics on this subscriber profile.

A maximum of one accounting policy can be associated with a profile at one time. Accounting policies are configured in the **config>log** context.

The **no** form of this command removes the accounting policy association.

Parameters *acct-policy-id* — Specifies the accounting *policy-id* as configured in the **config>log>accounting-policy** context.

Values 1 to 99

accu-stats-policy

Syntax **accu-stats-policy** *policy-name*
no accu-stats-policy

Context config>subscr-mgmt>sub-profile

Description This command associates an accumulated statistics policy with a subscriber profile.

The **no** form of this command removes the association of the accu-stats-policy from the subscriber profile. It is possible to remove the policy from the subscriber profile while the subscriber is still online, however, the statistics remain in memory and must be cleared manually, using the **clear subscriber-mgmt accu-stats active-subs no-accu-stats-policy** command.

Parameters *policy-name* — Specifies the name of the accumulated statistics policy to be used, to a maximum of 32 characters.

ancp-policy

Syntax **ancp-policy** *name*

Context config>subscr-mgmt>sub-prof>ancp

Description This command specifies an existing Access Node Control Protocol (ANCP) policy to associate with the subscriber profile. The policy is associated with either the ANCP string (static case) or subscriber-profile (dynamic case) and defines the behavior of the hosts belonging to these profiles.

The **no** form of this command removes the policy name from the ANCP configuration.

Parameters *name* — Specifies an existing ANCP policy name up to 32 characters.

collect-stats

Syntax [**no**] **collect-stats**

Context config>subscr-mgmt>sub-prof

Description When enabled, the agent collects non-RADIUS accounting statistics.

When the **no collect-stats** command is issued the statistics are still accumulated by the IOM cards. However, the CPU will not obtain the results and write them to the billing file. If a subsequent **collect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

Default collect-stats

agg-rate-limit

Syntax **agg-rate-limit** {**max** | *kilobits-per-second*} [**queue-frame-based-accounting**]
no agg-rate-limit

Context config>subscr-mgmt>sub-prof>egress

Description This command define a subscriber aggregate limit when the subscriber profile is directly associated with an egress port based scheduler instead of a scheduler policy. The optional queue-frame-based-accounting keyword allows the subscriber queues to operate in the frame based accounting mode.

Once egress frame based accounting is enabled on the subscriber profile, all queues associated with the subscriber (created through the sla-profile associated with each subscriber host) will have their rate and CIR values interpreted as frame based values. When shaping, the queues will include the 12-byte Inter-Frame Gap (IFG) and 8-byte preamble for each packet scheduled out the queue. The profiling CIR threshold will also include the 20-byte frame encapsulation overhead. Statistics associated with the queue do not include the frame encapsulation overhead. Packet byte offset settings are not included in the applied rate when queue frame based accounting is configured, however the offsets are applied to the statistics.

The **queue-frame-based-accounting** keyword does not change the behavior of the egress-agg-rate-limit rate value. Since the egress-agg-rate-limit is always associated with egress port based scheduling and egress port based scheduling is dependent on frame based operation, the egress-agg-rate-limit rate is always interpreted as a frame based value.

Enabling queue-frame-based-accounting will not cause statistics for queues associated with the subscriber to be cleared.

The **no** form of this command removes both an egress aggregate rate limit and egress frame based accounting for all subscribers associated with the sub-profile. If a subscriber's accounting mode is changed, the subscriber's queue statistics are cleared.

Parameters {**max |** *kilobits-per-second*} — Specifies the egress aggregate rate limit. The **max** keyword and *kilobits-per-second* parameter are mutually exclusive. Either **max** or a value for kilobits-per-second must follow the **egress-agg-rate-limit** command.

max — Specifies that the egress aggregate rate limit for the subscriber is unlimited. Scheduling for the subscriber queues will only be governed by the individual queue parameters and any congestion on the port relative to each queues scheduling priority.

kilobits-per-second — Specifies an actual egress aggregate rate to which all queues associated with the sub-profile is limited. The limit is managed per subscriber associated with the sub-profile. The value must be defined as an integer and is representative of increments of 1000 bits per second.

Values 1 to 800000000, max

Default max

queue-frame-based-accounting — Enables frame based accounting on all queues associated with the subscriber profile. If frame based accounting is required when a subscriber aggregate limit is not necessary, the max keyword should precede the queue-frame-based-accounting keyword. If frame based accounting must be disabled, execute the **egress-agg-rate-limit** command without the **queue-frame-based-accounting** keyword present.

Default Frame based accounting is disabled by default.

queue-frame-based-accounting — Specifies whether to use frame-based accounting when evaluating the aggregation rate limit for the egress queues for this SAP

avg-frame-size

Syntax **avg-frame-size** *bytes*
no avg-frame-size

Context config>subscr-mgmt>sub-profile>egress

Description This command specifies the average frame size used in the calculation of the fixed and variable encapsulation offset when the **encap-offset** command is enabled in the egress context of a subscriber profile.

If the user does not explicitly configure a value for the avg-frame-size parameter, then it will also be assumed the offset is zero.

The **no** form of this command removes the avg-frame-size parameter from the subscriber profile.

Default avg-frame-size 0

Parameters *bytes* — Specifies the average frame size value to be used in the adjustment of the subscriber aggregate rate to account for the per packet variable expansion of the last mile for the specific session used by the subscriber host.

Values 64 to 4096

encap-offset

Syntax **encap-offset** [**type** *type*]
no encap-offset

Context config>subscr-mgmt>sub-profile>egress

Description This command enables the adjustment of the queue and subscriber aggregate rate based on the last mile Ethernet or ATM encapsulation.

The data path computes the adjusted frame size real-time for each serviced packet from a queue by adding the actual packet size to the fixed offset provided by CPM for this queue and variable AAL5 padding.

When this command is enabled, the fixed packet offset is derived from the encapsulation type value signaled in the Access-loop-encapsulation sub-TLV in the Vendor-Specific PPPoE Tags or DHCP Relay Options as per RFC 4679. If the user specifies an encapsulation type with the command, this value is used as the default value for all hosts of this subscriber until a host session signaled a valid value. The signaled value is applied to this host only and the remaining hosts of this subscriber continue to use the user entered default type value if configured, or no offset is applied. However, hosts of the same subscriber using the same SLA profile and which are on the same SAP will share the same instance of FC queues. In this case, the last valid encapsulation value signaled by a host of that same instance of the SAP egress QoS policy will override any previous signaled or configured value.

If the user manually applied a constant byte offset to each packet serviced by the queue by configuring the packet-byte-offset, it will have no effect on the net offset computed for the packet. This net offset is stored in the subscriber host table.

The procedures for handling signaling changes or configuration changes affecting the subscriber profile are as follows:

The avg-frame-size parameter in the subscriber profile is ignored.

If the user specifies an encapsulation type with the command, this value is used as the default value for all hosts of this subscriber until a host session signaled a valid value. The signaled value is applied to this host and other hosts of the same subscriber sharing the same SLA profile and which are on the same SAP. The remaining hosts of this subscriber continue to use the user entered default type value if configured, or no offset is applied.

If the user enables/disables the encap-offset option, or changes the parameter value of the encap-offset option, CPM immediately triggers a re-evaluation of subscribers hosts using the corresponding subscriber profile and an update the IOM with the new fixed offset value.

If a subscriber has a static host or an ARP host, the subscriber host continues to use the user-configured default encapsulation type value or the last valid encapsulation value signaled in the PPPoE tags or DHCP relay options by other hosts of the same subscriber which use the same SLA profile instance. If none was signaled or configured, then no rate adjustment is applied.

When the encap-offset option is configured in the subscriber profile, the subscriber host queue rates, that is, CLI and operational PIR and CIR as well as queue bucket updates, the queue statistics, that is, forwarded, dropped, and HQoS offered counters use the last-mile frame-over-the-wire format. The scheduler policy CLI and operational rates also use LM-FoW format. The port scheduler max-rate and the priority level rates and weights, if a Weighted Scheduler Group is used, are always entered in CLI and interpreted as local port frame-over-the-wire rates. The same is true for an agg-rate-limit applied to a Vport. Finally the subscriber agg-rate-limit is entered in CLI as last-mile frame-over-the-wire rate. The system maintains a running average frame expansion ratio for each queue to convert queue rates between these two formats.

The **no** form of this command reverts to the default.

Parameters *type* — The name of the default encapsulation used for all host queues of a subscriber in the absence of a valid value signaled in the PPPoE tags.

Values pppoa-llc, pppoa-null, pppoeoa-llc, pppoeoa-llc-fcs, pppoeoa-llc-tagged, pppoeoa-llc-tagged-fcs, pppoeoa-null, pppoeoa-null-fcs, pppoeoa-null-tagged, pppoeoa-null-tagged-fcs, ipoa-llc, ipoa-null, ipoeoa-llc, ipoeoa-llc-fcs, ipoeoa-llc-tagged, ipoeoa-llc-tagged-fcs, ipoeoa-null, ipoeoa-null-fcs, ipoeoa-null-tagged, ipoeoa-null-tagged-fcs, pppoe, pppoe-tagged, ipoe, ipoe-tagged

scheduler

Syntax **scheduler** *scheduler-name* **rate** *pir-rate* [**cir** *cir-rate*]
no scheduler *scheduler-name*

Context config>subscr-mgmt>sub-prof>egress>sched
config>subscr-mgmt>sub-prof>ingress>sched

Description This command provides a way to override parameters of the existing scheduler associated with the egress or ingress scheduler policy. A scheduler defines bandwidth controls that limit each child (other schedulers and queues) associated with the scheduler. Scheduler objects are created within the hierarchical tiers of the policy. It is assumed that each scheduler created will have queues or other schedulers defined as child associations. The scheduler can be a child (take bandwidth from a scheduler in a higher tier).

The **no** form of this command reverts to the default.

Parameters *scheduler-policy-name* — Specify an existing scheduler policy name.

pir-rate — Specify the pir-rate, in kilobits, to override the administrative PIR used by the scheduler. When the **rate** command is executed, a valid PIR setting must be explicitly defined
Fractional values are not allowed and must be given as a positive integer.

Values 1 to 3200000000, max

cir-rate — The **cir** parameter overrides the administrative CIR used by the scheduler. When the **rate** command is executed, a CIR setting is optional. The sum keyword specifies that the CIR be used as the summed CIR values of the children schedulers or queues.
Fractional values are not allowed and must be given as a positive integer.

Values 0 to 3200000000, sum, max

Default sum

scheduler-policy

Syntax **scheduler-policy** *scheduler-policy-name*
no scheduler-policy

Context config>subscr-mgmt>sub-profile>egress
config>subscr-mgmt>sub-profile>ingress

Description This command specifies a scheduler policy to associate to the subscriber profile. Scheduler policies are configured in the **configure>qos>scheduler>policy** context. Each scheduler policy is divided up into groups of schedulers based on the tier each scheduler is created under. A tier is used to give structure to the schedulers within a policy and define rules for parent scheduler associations. The policy defines the hierarchy and operating parameters for virtual schedulers.

The **no** form of this command reverts to the default.

Parameters *scheduler-policy-name* — Specify an existing scheduler policy name up to 32 characters.

hs-agg-rate-limit

Syntax **hs-agg-rate-limit** *kilobits-per-second*
no hs-agg-rate-limit

Context config>subscr-mgmt>sub-profile>egress

Description This command configures the aggregate rate limit for the subscriber in single SLA mode. In single SLA mode, the **hs-aggregate-rate-limit** command under the **config>subscr-mgmt>sla-profile** context should not be configured.

When both, **sub-profile>egress>hs-agg-rate-limit** and **sla-profile>egress>hs-agg-rate-limit** are configured, the system takes the minimum of the two to program the subscriber's aggregate rate.

The **no** form of this command removes the value from the configuration.

Parameters *kilobits-per-second* — Specifies the HS egress aggregate rate limit.

Values 1 to 100000000

lag-per-link-hash

Syntax **lag-per-link-hash class** {*class*} **weight** *weight*
no lag-per-link-hash

Context config>subscr-mgmt>sub-profile>egress

Description This command configures weight and class to be used on LAG egress when the LAG uses weighted per-link-hash by subscribers with the profile assigned. Subscribers using profile with lag-per-link-hash default configuration, inherit weight and class from the SAP configuration (1 and 1 respectively if none configured under SAP).

The **no** form of this command restores default configuration.

Parameters *class* — Specifies the class to be used to select a LAG link.

Values 1, 2, 3

Default 1

weight — Specifies the weight to be associated with this SAP when selecting a LAG link.

Values 1 to 1024

Default 1

policer-control-policy

Syntax **policer-control-policy** *policy-name* [**create**]
no policer-control-policy

Context config>subscr-mgmt>sub-prof>ingress
config>subscr-mgmt>sub-prof>egress

Description This command is used to create, delete, or modify policer control policies. The **policer-control-policy** controls the aggregate bandwidth available to a set of child policers. Once created, the policy can be applied to ingress or egress SAPs. The policy can also be applied to the ingress or egress context of a sub-profile.

The **no** form of this command reverts to the default.

Parameters *policy-name* — Specifies the policer control policy name. Each policer-control-policy must be created with a unique policy name. The name given as policy-name must adhere to the system policy ASCII naming requirements. If the defined policy-name already exists, the system will enter that policy's context for editing purposes. If policy-name does not exist, the system will attempt to create a policy with the specified name. Creating a policy may require use of the create parameter when the system is configured for explicit object creation mode.

create — The **create** keyword is required when a new policy is being created and the system is configured for explicit object creation mode.

max-rate

Syntax **max-rate** {*rate* | **max**}
no max-rate

Context config>subscr-mgmt>sub-prof>ingress>policer-control-policy
config>subscr-mgmt>sub-prof>egress>policer-control-policy

Description The **max-rate** command defines the parent policer's PIR leaky bucket's decrement rate. A parent policer is created for each time the policer-control-policy is applied to either a SAP or subscriber instance. Packets that are not discarded by the child policers associated with the SAP or subscriber instance are evaluated against the parent policer's PIR leaky bucket.

For each packet, the bucket is first decremented by the correct amount based on the decrement rate to derive the current bucket depth. The current depth is then compared to one of two discard thresholds associated with the packet. The first discard threshold (discard-unfair) is applied if the FIR (Fair Information Rate) leaky bucket in the packet's child policer is in the confirming state. The second discard threshold (discard-all) is applied if the child policer's FIR leaky bucket is in the exceed state. Only one of the two thresholds is applied per packet. If the current depth of the parent policer PIR bucket is less than the threshold value, the parent PIR bucket is in the conform state for that particular packet. If the depth is equal to or greater than the applied threshold, the bucket is in the violate state for the packet.

If the result is "conform," the bucket depth is increased by the size of the packet (plus or minus the per-packet-offset setting in the child policer) and the packet is not discarded by the parent policer. If the result is "violate," the bucket depth is not increased and the packet is discarded by the parent policer. When the parent policer discards a packet, any bucket depth increases (PIR, CIR and FIR) in the parent policer caused by the packet are canceled. This prevents packets that are discarded by the parent policer from consuming the child policers PIR, CIR and FIR bandwidth.

The **policer-control-policy root max-rate** setting may be overridden on each SAP or sub-profile where the policy is applied.

The **no** form of this command reverts to the default.

Default max-rate max

Parameters *rate* — Specifies the max rate in kb/s. Defining the *kilobits-per-second* value is mutually exclusive with the **max** parameter. The *kilobits-per-second* value must be defined as an integer that represents the number of kilobytes that the parent policer is decremented per second. The actual decrement is performed per packet based on the time that has elapsed since the last packet associated with the parent policer.

Values 0 – 2000000000

max — Specifies the maximum frame-based bandwidth limit of this policer. The **max** parameter is mutually exclusive with defining a **kilobits-per-second** value. When max is specified, the parent policer does not enforce a maximum rate on the aggregate throughput of the child policers. This is the default setting when the **policer-control-policy** is first created and is the value that the parent policer returns to when no max-rate is executed. In order for the parent policer to be effective, a kilobits-per-second value should be specified.

priority-mbs-thresholds

Syntax **priority-mbs-thresholds**

Context config>subscr-mgmt>sub-prof>ingress>policer-control-policy
config>subscr-mgmt>sub-prof>egress>policer-control-policy

Description The **priority-mbs-thresholds** command contains the root arbiter parent policer's **min-thresh-separation** command and each priority level's **mbs-contribution** command that is used to internally derive each priority level's shared-portion and fair-portion values. The system uses each priority level's shared-portion and fair-portion value to calculate each priority level's discard-unfair and discard-all MBS thresholds that enforce priority sensitive rate-based discards within the root arbiter's parent policer.

The **priority-mbs-thresholds** CLI node always exists and does not need to be created.

min-thresh-separation

Syntax **min-thresh-separation** *size* [**bytes** | **kilobytes**]
no min-thresh-separation

Context config>subscr-mgmt>sub-prof>ingress>policer-control-policy>priority-mbs-thresholds
config>subscr-mgmt>sub-prof>egress>policer-control-policy>priority-mbs-thresholds

Description The **min-thresh-separation** command defines the minimum required separation between each in-use discard threshold maintained for each parent policer context associated with the policer-control-policy. The min-thresh-separation value may be overridden on each SAP or sub-profile to which the policy is applied.

The system uses the default or specified min-thresh-separation value in order to determine the minimum separation required between each of the of the parent policer discard thresholds. The system enforces the minimum separation based on the following behavior in two ways. The first is determining the size of the shared-portion for each priority level (when the **mbs-contribution** command's optional fixed keyword is not specified):

- When a parent policer instance's priority level has less than two child policers associated, the shared-portion for the level is zero.
- When a parent policer instance's priority level has two or more child policers associated, the shared-portion for the level is equal to the current value of **min-thresh-separation**.

The second function the system uses the **min-thresh-separation** value for is determining the value per priority level for the fair-portion:

- When a parent policer instance's priority level has no child policers associated, the fair-portion for the level is zero.
- When a parent policer instance's priority level has one child policer associated, the fair-portion is equal to the maximum of the min-thresh-separation value and the priority level's mbs-contribution value.
- When a parent policer instance's priority level has two or more child policers associated, the fair-portion is equal to the maximum of the following:
  - **min-thresh-separation** value
  - The priority level's **mbs-contribution** value less **min-thresh-separation** value

When the **mbs-contribution** command's optional fixed keyword is defined for a priority level within the policy, the system will treat the defined **mbs-contribution** value as an explicit definition of the priority level's MBS. While the system will continue to track child policer associations with the parent policer priority levels, the association counters will have no effect. Instead the following rules is used to determine a fixed priority level's shared-portion and fair-portion:

- If a fixed priority level's **mbs-contribution** value is set to zero, both the shared-portion and fair-portion is set to zero
- If the **mbs-contribution** value is not set to zero:
  - The shared-portion is set to the current **min-thresh-separation** value
  - The fair-portion is set to the maximum of the following:

min-thresh-separation value

mbs-contribution value less **min-thresh-separation value**

Each time the **min-thresh-separation** value is modified, the thresholds for all instances of the parent policer created through association with this **policer-control-policy** are reevaluated.

Determining the Correct Value for the Minimum Threshold Separation Value

The minimum value for **min-thresh-separation** should be set equal to the maximum size packet that is handled by the parent policer. This ensures that when a lower priority packet is incrementing the bucket, the size of the increment will not cause the bucket's depth to equal or exceed a higher priority threshold. It also ensures that an unfair packet within a priority level cannot cause the PIR bucket to increment to the discard-all threshold within the priority.

When evaluating maximum packet size, each child policer's per-packet-offset setting should be taken into consideration. If the maximum size packet is 1518 bytes and a per-packet-offset parameter is configured to add 20 bytes per packet, min-thresh-separation should be set to 1538 due to the fact that the parent policer will increment its PIR bucket using the extra 20 bytes.

In most circumstances, a value larger than the maximum packet size is not necessary. Management of priority level aggregate burst tolerance is intended to be implemented using the priority level **mbs-contribution** command. Setting a value larger than the maximum packet size will not adversely affect the policer performance, but it may increase the aggregate burst tolerance for each priority level.

Note: A priority level's shared-portion of the parent policer's PIR bucket depth is only necessary to provide some separation between a lower priority's discard-all threshold and this priority's discard-unfair threshold. It is expected that the burst tolerance for the unfair packets is relatively minimal since the child policers feeding the parent policer priority level all have some amount of fair burst before entering into an FIR exceed or unfair state. The fair burst amount for a priority level is defined using the mbs-contribution command.

The **no** form of this command returns the policy's **min-thresh-separation** value to the default value.

Parameters *size* — The size parameter is required when executing the **min-thresh-separation** command. It is expressed as an integer and specifies the shared portion in bytes or kilobytes that is selected by the trailing bytes or kilobytes keywords. If both bytes and kilobytes are missing, kilobytes is the assumed value. Setting this value has no effect on parent policer instances where the **min-thresh-separation** value has been overridden.

Values 0 to 16777216

[**bytes** | **kilobytes**] — The **bytes** keyword is optional and is mutually exclusive with the **kilobytes** keyword. When specified, size is interpreted as specifying the size of **min-thresh-separation** in bytes.

The **kilobytes** keyword is optional and is mutually exclusive with the **bytes** keyword. When specified, size is interpreted as specifying the size of **min-thresh-separation** in kilobytes.

Default kilobytes

priority

Syntax **priority** *level*

Context config>subscr-mgmt>sub-prof>ingress>policer-control-policy>priority-mbs-thresholds
config>subscr-mgmt>sub-prof>egress>policer-control-policy>priority-mbs-thresholds

Description The **priority** level command contains the **mbs-contribution** configuration command for a given strict priority level. Eight levels are supported numbered 1 through 8 with 8 being the highest strict priority.

Each of the eight priority CLI nodes always exists and do not need to be created. While parameters exist for each priority level, the parameters are only applied when the priority level within a parent policer instance is currently supporting child policers.

Parameters *level* — Specifies the priority level override.

Values 1 to 8

mbs-contribution

Syntax **mbs-contribution** *size* [**bytes** | **kilobytes**] [**fixed**]
no mbs-contribution

Context config>subscr-mgmt>sub-prof>ingress>policer-control-policy>priority-mbs-thresholds>priority
config>subscr-mgmt>sub-prof>egress>policer-control-policy>priority-mbs-thresholds>priority

Description The **mbs-contribution** command is used to configure the policy-based burst tolerance for a parent policer instance created when the policy is applied to a SAP or subscriber context. The system uses the parent policer's **min-thresh-separation** value, the priority level's **mbs-contribution** value and the number of child policers currently attached to the priority level to derive the priority level's shared-portion and fair-portion of burst tolerance within the local priority level. The shared-portion and fair-portions for each priority level are then used by the system to calculate each priority level's discard-unfair threshold and discard-all threshold.

The value for a priority level's **mbs-contribution** within the policer-control-policy may be overridden on the SAP or subscriber sub-profile where the policy is applied in order to allow fine tuning of the discard-unfair and discard-all thresholds relevant to the needs of the local child policers on the object.

Accumulative Nature of Burst Tolerance for a Parent Policer Priority Level

When defining **mbs-contribution**, the specified size may only be a portion of the burst tolerance associated with the priority level. The packets associated with the priority level share the burst tolerance of lower within the parent policer. As the parent policer PIR bucket depth increases during congestion, the lower priority packets eventually experience discard based on each priority's discard-unfair and discard-all thresholds. Assuming congestion continues once all the lower priority packets have been prevented from consuming bucket depth, the burst tolerance for the priority level is consumed by its own packets and any packets associated with higher priorities.

The Effect of Fair and Unfair Child Policer Traffic at a Parent Policer Priority Level

The system continually monitors the offered rate of each child policer on each parent policer priority level and detects when the policer is in a congested state (the aggregate offered load is greater than the decrement rate defined on the parent policer). As previously stated, the result of congestion is that the parent policer's bucket depth will increase until it eventually hovers around either a discard-unfair or discard-all threshold belonging to one of the priority levels. This threshold is the point where enough packets are being discarded that the increment rate and decrement rate begin to even out. If only a single child policer is associated to the priority level, the discard-unfair threshold is not used since fairness is only applicable when multiple child policers are competing at the same priority level.

When multiple child policers are sharing the congested priority level, the system uses the offered rates and the parenting parameters of each child to determine the fair rate per child when the parent policer is unable to meet the bandwidth needs of each child. The fair rate represents the amount of bandwidth that each child at the priority level should receive relative to the other children at the same level according to the policer control policy instance managing the child policers. This fair rate is applied as the decrement rate for each child's FIR bucket. Changing a child's FIR rate does not modify the amount of packets forwarded by the parent policer for the child's priority level. It simply modifies the forwarded ratio between the children on that priority level. Since each child FIR bucket has some level of burst tolerance before marking its packets as unfair, the current parent policer bucket depth may at times rise above the discard-unfair threshold. The mbs-contribution value provides a means to define how much separation is provided between the priority level's discard-unfair and discard-all threshold to allow the parent policer to absorb some amount of FIR burst before reaching the priority's discard-all threshold.

This level of fair aggregate burst tolerance is based on the decrement rate of the parent policer's PIR bucket while the individual fair bursts making up the aggregate are based on each child's FIR decrement rate. The aggregate fair rate of the priority level is managed by the system with consideration of the current rate of traffic in higher priority levels. In essence, the system ensures that for each iteration of the child FIR rate calculation, the sum of the child FIR decrement rates plus the sum of the higher priority traffic increment rates equals the parent policers decrement rate. This means that dynamic amounts of higher priority traffic can be ignored when sizing a lower priority's fair aggregate burst tolerance. Consider the following:

- The parent policer decrement rate is set to 20 Mb/s (max-rate 20,000).
- A priority level's fair burst size is set to 30 kbytes (mbs-contribution 30 kbytes).
- Higher priority traffic is currently taking 12 Mb/s.

- The priority level has three child policers attached.
- Each child's PIR MBS is set to 10 kbytes, which makes each child's FIR MBS 10 kbytes.
- The children want 10 Mb/s, but only 8 Mb/s is available,
- Based on weights, the children's FIR rates are set as follows:

	FIR Rate	**FIR MBS**
Child 1	4 Mb/s	10 kbytes
Child 2	3 Mb/s	10 kbytes
Child 3	1 Mb/s	10 kbytes

The 12 Mb/s of the higher priority traffic and the 8 Mb/s of fair traffic equal the 20 Mb/s decrement rate of the parent policer.

It is clear that the higher priority traffic is consuming 12 Mb/s of the parent policer's decrement rate, leaving 8 Mb/s of decrement rate for the lower priority's fair traffic.

- The burst tolerance of child 1 is based on 10 kbytes above 4 Mb/s,
- The burst tolerance of child 2 is based on 10 kbytes above 3 Mb/s,
- The burst tolerance of child 3 is based on 10 kbytes above 1 Mb/s.

If all three children burst simultaneously (unlikely), they will consume 30 kbytes above 8 Mb/s. This is the same as the remaining decrement rate after the higher priority traffic.

Parent Policer Total Burst Tolerance and Downstream Buffering

The highest in-use priority level's discard-all threshold is the total burst tolerance of the parent policer. In some cases the parent policer represents downstream bandwidth capacity and the max-rate of the parent policer is set to prevent overrunning the downstream bandwidth. The burst tolerance of the parent policer defines how much more traffic may be sent beyond the downstream scheduling capacity. In the worst case scenario, when the downstream buffering is insufficient to handle the total possible burst from the parent policer, downstream discards based on lack of buffering may occur. However, in all likelihood, this is not the case.

In most cases, lower priority traffic in the policer is responsible for the greater part of congestion above the parent policer rate. Since this traffic is discarded with a lower threshold, this lowers the effective burst tolerance even while the highest priority traffic is present.

Configuring a Priority Level's MBS Contribution Value

In the most conservative case, a priority level's **mbs-contribution** value may be set to be greater than the sum of child policer's mbs and one max-size-frame per child policer. This ensures that even in the absolute worst case where all the lower priority levels are simultaneously bursting to the maximum capacity of each child, enough burst tolerance for the priority's children will exist if they also burst to their maximum capacity.

Since simply adding up all the child policer's PIR MBS values may result in large overall burst tolerances that are not ever likely to be needed, you should consider some level of burst oversubscription when configuring the **mbs-contribution** value for each priority level. The amount of oversubscription should be determined based on the needs of each priority level.

Using the Fixed Keyword to Create Deterministic Parent Policer Discard Thresholds

In the default behavior, the system ignores the **mbs-contribution** values for a priority level on a subscriber or SAP parent policer when a child policer is not currently associated with the level. This prevents additional burst tolerance from being added to higher priority traffic within the parent policer.

This does cause fluctuations in the defined threshold values when child policers are added or removed from a parent policer instance. If this behavior is undesirable, the fixed keyword may be used which causes the **mbs-contribution** value to always be included in the calculation of parent policer's discard thresholds. The defined **mbs-contribution** value may be overridden on a subscriber sla-profile or on a SAP instance, but the fixed nature of the contribution cannot be overridden.

If the defined **mbs-contribution** value for the priority level is zero, the priority level will have no effect on the parent policer's defined discard thresholds. A packet associated with the priority level will use the next lower priority level's discard-unfair and discard-all thresholds.

The **no mbs-contribution** command returns the policy's priority level's MBS contribution to the default value. When changed, the thresholds for the priority level and all higher priority levels for all instances of the parent policer is recalculated.

Parameters *size* [**bytes** | **kilobytes**] — The size parameter is required when executing the **mbs-contribution** command. It is expressed as an integer and specifies the priority's specific portion amount of accumulative MBS for the priority level in bytes or kilobytes which is selected by the trailing **bytes** or **kilobytes** keywords. If both **bytes** and **kilobytes** are missing, **kilobytes** is assumed. Setting this value has no effect on parent policer instances where the priority level's **mbs-contribution** value has been overridden.

Values 0 to 16777216

bytes | **kilobytes**: — The **bytes** keyword is optional and is mutually exclusive with the **kilobytes** keyword. When specified, size is interpreted as specifying the size of **min-thresh-separation** in bytes.

The **kilobytes** keyword is optional and is mutually exclusive with the **bytes** keyword. Size is interpreted as specifying the size of min-thresh-separation in kilobytes.

Default **kilobytes**

fixed — This optional fixed keyword is used to force the inclusion of the defined **mbs-contribution** value in the parent policer's discard threshold calculations. If the **mbs-contribution** command is executed without the **fixed** keyword, the fixed calculation behavior for the priority level is removed.

radius-accounting

Syntax **radius-accounting**

Context config>subscr-mgmt>sub-prof

Description This command enters the RADIUS accounting context for configuring RADIUS accounting subscriber profile parameters.

policy

Syntax **policy** *name1* [*name2*] [*name3*] [*name4*] [*name5*]
no policy

Context config>subscr-mgmt>sub-prof>rad-acct

Description This command specifies the RADIUS accounting policy for the subscriber that is using this subscriber profile. This command allows the configuration of up to five RADIUS accounting policies. The RADIUS accounting policies function according to their respective configuration, including the individual accounting mode, their own included attributes, and the update interval.

Parameters *name1* — Specifies the name of the RADIUS accounting policy to be used for the subscriber profile (32 characters maximum).

name2 — Specifies the name of the second RADIUS accounting policy to be used for the subscriber profile (32 characters maximum).

name3 — Specifies the name of the third RADIUS accounting policy to be used for the subscriber profile (32 characters maximum).

name4 — Specifies the name of the fourth RADIUS accounting policy to be used for the subscriber profile (32 characters maximum).

name5 — Specifies the name of the fifth RADIUS accounting policy to be used for the subscriber profile (32 characters maximum).

session-optimized-stop

Syntax [**no**] **session-optimized-stop**

Context config>subscr-mgmt>sub-prof>rad-acct

Description This command optimizes a RADIUS Accounting Stop message for a PPPoE session termination (specifically for session accounting mode when the host update is enabled). By default when a PPPoE session terminates, the system removes a dual stack host in sequence, one host at a time. Therefore, the system will generate a RADIUS accounting interim for each host termination until only the final host is left. The final host will generate a final accounting stop message. Enabling this command will trigger a single Stop RADIUS accounting message and include information of all hosts without the host updates.

sla-profile-map

Syntax **sla-profile-map**

Context config>subscr-mgmt>sub-prof

Description This command enables the context to configure SLA profile mapping.

entry

Syntax **entry key** *sub-profile-string* **sub-profile** *sub-profile-name*
no entry key *sub-profile-string*

Context config>subscr-mgmt>sub-prof>sla-prof-map

Description This command configures SLA profile string mappings.

Parameters *sub-profile-string* — Specifies the subscriber profile string up to 16 characters.

sub-profile-name — Specifies the subscriber profile name up to 32 characters.

use-direct-map-as-default

Syntax [**no**] **use-direct-map-as-default**

Context config>subscr-mgmt>sub-prof>sla-prof-map

Description This command enables direct mapping of the SLA profile as default.

The **no** form of this command disables direct mapping.

sub-mcac-policy

Syntax **sub-mcac-policy** *policy-name*
no sub-mcac-policy

Context config>subscr-mgmt>sub-prof

Description This command references the policy template in which the mcac bandwidth limits are defined. Mcac for the subscriber is effectively enabled with this command when the sub-profile is applied to the subscriber. The bandwidth of the channels is defined in a different policy (under the **config>router>mcac** context) and this policy is applied on the interface level as follows:

- For group-interfaces under the **config>service>vprn>igmp>grp-if>mcac** context
- For regular interfaces under the **config>service/router>igmp>interface>mcac** context

In case of HQoS Adjustment, it is mandatory that the sub-mcac-policy be created and applied to the subscriber. The sub-mac-policy does not have to contain any bandwidth constrains, but it has to be in a no shutdown state in order for HQoS Adjustment to work.

The **no** form of this command removes the policy from the configuration.

Parameters *policy-name* — Specifies the policy name configured in the config>subscr-mgmt>sub-mcac-policy context.

secondary-shaper-hashing

Syntax [**no**] **secondary-shaper-hashing**

Context config>subscr-mgmt>sub-prof

Description This command enables LAG secondary shaper ID hashing (the equivalent of Vport hashing) for the HSMDA. With this feature enabled, secondary shaper ID hashing can span multiple forwarding complexes on egress LAG. The default is to perform secondary shaper ID hashing on egress and requires all active LAG members to be on the same forwarding complex.

The **no** form of this command enables the default behavior.

vport-hashing

Syntax [**no**] **vport-hashing**

Context config>subscr-mgmt>sub-prof

Description This command enables LAG Vport ID hashing. When enabled, Vport ID hashing can span multiple forwarding complexes on egress LAG. The default is to perform Vport ID hashing on egress and requires all active LAG members to be on the same forwarding complex. This command is not applicable to the HSMDA.

Note: LAG hashing parameters that are configured under **config>lag**, for example, **per-link-hash**, take precedence and are incompatible with the **vport-hashing** command.

The **no** form of this command enables the default behavior.

volume-stats-type

Syntax **volume-stats-type** {**ip** | **default**}
no volume-stats-type

Context config>subscr-mgmt>sub-prof

Description This command enables the reporting of Layer 3 (IP) based subscriber host volume accounting data.

By default, subscriber host volume accounting data includes Layer 2 header octets and can be configured to include a fixed packet byte offset or last-mile encapsulation overhead.

The **no** form of this command reverts to the default.

Default volume-stats-type default

Parameters **default** — Specifies that the subscriber host volume accounting data is reported including the Layer 2 header octets and optional delta's introduced by configuration (for example: packet byte offset, last mile aware shaping, and so on).

ip — Specifies that the subscriber host volume accounting data reporting is based on Layer 3 (IP) packet sizes. This includes subscriber host ingress/egress queue and policer stats in snmp, CLI show commands, RADIUS and XML accounting, and Diameter Gx usage monitoring. RADIUS and Diameter (DCCA) based credit control volume quota are interpreted as Layer 3 (IP).

igmp-policy

Syntax **igmp-policy** *policy-name*
no igmp-policy

Context config>subscr-mgmt>sub-prof

Description This command will enable IGMP processing per subscriber host. Without this command IGMP states will not be maintained per subscriber hosts. The referenced policy is defined under the **configure>subscr-mgmt** context and can be only applied via the sub-profile.

The referenced policy contains entries such as:

- description statement
- import statement — IGMP filters
- egress-rate-modify statement—HQoS Adjustment
- mcast-redirection statement—redirection to alternate interface
- static statement—definition of static IGMP groups
- version statement —IGMP version
- fast-leave statement
- max-num-groups statement—the maximum number of multicast groups allowed

The **no** form of this command reverts to the default.

Parameters *policy-name* — Specifies the name of the IGMP policy for the subscriber. The policy itself is defined under the **configure>sub-mgmt** context.

hs-sla-mode

Syntax **hs-sla-mode** {**expanded** | **single**}
no hs-sla-mode

Context config>subscr-mgmt>sub-prof

Description This command specifies the SLA profile handling mode for the subscriber if on an HS board.

The **no** form of this command reverts to the default.

Parameters **expanded** — Specifies the expanded SLA profile handling mode for the subscriber if on an HS board.

single — Specifies a single SLA profile handling mode for the subscriber if on an HS board.

hsmda

Syntax **hsmda**

Context config>subscr-mgmt>sub-prof

Description This command enables the context to configure egress and ingress HSMDA queue parameters.

egress-qos

Syntax **egress-queues**

Context config>subscr-mgmt>sub-prof>hsmda

Description This command enables the context to configure SAP egress QOS policy for the HSMDA egress queue.

ingress-qos

Syntax **ingress-queues**

Context config>subscr-mgmt>sub-prof>hsmda>egress-queues

Description This command enables the context to configure SAP egress QOS policy for the HSMDA ingress queue

agg-rate

Syntax **agg-rate** *rate*
no agg-rate

Context config>port>sonet-sdh>path>access>egress>vport
config>port>ethernet>access>egress>vport

Description This command configures an aggregate rate for the Vport. The **agg-rate** *rate*, **port-scheduler-policy** and **scheduler-policy** commands are mutually exclusive. Changing between the use of a scheduler policy and the use of an agg-rate/port-scheduler-policy involves removing the existing command and applying the new command. Applying a **scheduler-policy** to a Vport is only applicable to Ethernet interfaces.

The **no** form of this command reverts to the default.

Parameters *rate* — Specifies the rate limit for the Vport.

Values 1 to 800000000, max

limit-unused-bandwidth

Syntax [**no**] **limit-unused-bandwidth**

Context config>port>sonet-sdh>path>access>egress>vport
config>port>ethernet>access>egress>vport

Description This command specifies to limit the unused bandwidth and allow a tighter control in allocation of bandwidth by HQoS. When enabled, HQoS algorithm distributes any unused aggregate bandwidth between queues operating below their fair share rates. This allows a simplified aggregate rate protection while allocating bandwidth by HQoS.

The **no** form of this command reverts to the default.

agg-rate-limit

Syntax **agg-rate-limit** *agg-rate*
no agg-rate-limit

Context config>subscr-mgmt>sub-prof>hsmda>egress-qos

Description This command defines a maximum total rate for all subscriber egress queues for each subscriber associated with the sub-profile. The **egress-agg-rate-limit** command is mutually exclusive with the egress-scheduler-policy. When an egress-scheduler-policy is defined on the sub-profile, the **egress-agg-rate-limit** command will fail. If the egress-**agg-rate-limit** command is specified, at attempt to bind an **egress-scheduler-policy** to the sub-profile will fail.

A port scheduler policy must be applied on the egress port or channel the subscriber instance is bound to in order for the defined **egress-agg-rate-limit** command to take effect. The egress port scheduler enforces the aggregate queue rate as it distributes its bandwidth at the various port priority levels. The port scheduler stops offering bandwidth to member queues once it has detected that the aggregate rate limit has been reached.

If a port scheduler is not defined on the egress port, the queues are allowed to operate based on their own bandwidth parameters.

The **no** form of this command removes the aggregate rate limit from the sub-profile.

Parameters *agg-rate* — Defines the maximum aggregate rate the egress queues associated with the subscriber profile may operate. The value is specified in kilobits per second in a base 10 context. A value of 1 indicates a rate of 1000 bits per second.

Values 1 to 40000000, max kb/s

qos

Syntax **qos** *policy-id*
no qos

Context config>subscr-mgmt>sub-prof>hsmda>egress-qos

Description This command assigns a SAP egress QoS policy to the HSMDA egress queue.

The **no** form of this command removes the policy from the configuration.

Default qos 1

Parameters *policy-id* — Specifies the policy ID of an existing QoS SAP egress policy.

Values 1 to 65535

qos

Syntax **qos** *policy-id*
no qos

Context config>subscr-mgmt>sub-prof>hsmda>ingress-qos

Description This command assigns a SAP ingress QOS policy to the HSMDA ingress queue.

The **no** form of this command removes the policy from the configuration.

Default qos 1

Parameters *policy-id* — Specifies the policy ID of an existing QoS SAP egress policy.

Values 1 to 65535

packet-byte-offset

Syntax **packet-byte-offset** {**add** *add-bytes* | **subtract** *sub-bytes*}
no packet-byte-offset

Context config>subscr-mgmt>sub-prof>hsmda>egress-qos

Description This command adds or subtracts the specified number of bytes to the accounting function for each packet handled by the HSMDA queue. Normally, the accounting and leaky bucket functions are based on the Ethernet DLC header, payload and the 4-byte CRC (everything except the preamble and inter-frame gap). As an example, the **packet-byte-offset** command can be used to add the frame encapsulation overhead (20 bytes) to the queues accounting functions.

The accounting functions affected include:

- Offered High Priority / In-Profile Octet Counter
- Offered Low Priority / Out-of-Profile Octet Counter
- Discarded High Priority / In-Profile Octet Counter
- Discarded Low Priority / Out-of-Profile Octet Counter
- Forwarded In-Profile Octet Counter
- Forwarded Out-of-Profile Octet Counter
- Peak Information Rate (PIR) Leaky Bucket Updates
- Committed Information Rate (CIR) Leaky Bucket Updates
- Queue Group Aggregate Rate Limit Leaky Bucket Updates

The secondary shaper leaky bucket, scheduler priority level leaky bucket and the port maximum rate updates are not affected by the configured packet-byte-offset. Each of these accounting functions are frame based and always include the preamble, DLC header, payload and the CRC regardless of the configured byte offset.

The packet-byte-offset command accepts either add or subtract as valid keywords which define whether bytes are being added or removed from each packet traversing the queue. An example use case for subtracting bytes from each packet is an IP based accounting function. Given a Dot1Q encapsulation, the command **packet-byte-offset** subtract 14 would remove the DLC header and the Dot1Q header from the size of each packet for accounting functions only. The 14 bytes are not actually removed from the packet, only the accounting size of the packet is affected.

As inferred above, the variable accounting size offered by the **packet-byte-offset** command is targeted at the queue and queue group level. The **packet-byte-offset**, when set, applies to all queues in the queue group. The accounting size of the packet is ignored by the secondary shapers, the scheduling priority level shapers and the scheduler maximum rate. The actual on-the-wire frame size is used for these functions to allow an accurate representation of the behavior of the subscriber’s packets on an Ethernet aggregation network.

The **packet-byte-offset** value may be overridden at the queue-group level.

The **no** form of this command reverts to the default.

Parameters **add** *add-bytes* — Indicates that the byte value should be added to the packet for queue and queue group level accounting functions. Either the **add** or **subtract** keyword must be specified. The corresponding byte value must be specified when executing the packet-byte-offset command. The **add** keyword is mutually exclusive with the **subtract** keyword.

Values 0 to 31

subtract *sub-bytes* — Indicates that the byte value should be subtracted from the packet for queue and queue group level accounting functions. The **subtract** keyword is mutually exclusive with the **add** keyword. Either the **add** or **subtract** keyword must be specified. The corresponding byte value must be specified when executing the packet-byte-offset command.

Note: The minimum resulting packet size used by the system is 1 byte.

Values 1 to 64

queue

Syntax **queue** *queue-id* [**create**]
no queue *queue-id*

Context config>subscr-mgmt>sub-prof>hsmda>ingress-qos>qos

Description This command specifies the HSMDA queue mapping for all packets in point-to-point services and unicast destined packets in multipoint services. Point-to-point services include epipe and other VLL type services. Multipoint services include IES, VPLS and VPRN services. The **queue** command does not apply to multicast, broadcast or unknown unicast packets within multipoint services (the multicast, broadcast and unknown commands must be used to define the queue mapping for non-unicast packets within a forwarding class). For Epipe services, the **queue** *queue-id* mapping applies to all packets, regardless of the packets destination MAC address.

Each forwarding class has a default queue ID based on the intrinsic hierarchy between the forwarding classes. Executing the queue command within the HSMDA context of a forwarding class with a different queue ID than the default overrides the default mapping. Multiple forwarding classes may be mapped to the same HSMDA queue ID.

The **no** form of this command returns the HSMDA queue mapping for queue to the default mapping for the forwarding class.

Parameters *queue-id* — Specifies the queue ID to override.

Values 1 to 8

create — This keyword is mandatory while creating a new queue override.

rate

Syntax **rate** *pir-rate* [**cir** *cir-rate*]
no rate

Context config>subscr-mgmt>sub-prof>hsmda>egress-qos>qos>queue

Description This command can be used to override specific attributes of the specified queue's Peak Information Rate (PIR) and the Committed Information Rate (CIR) parameters. The PIR defines the maximum rate that the queue can transmit packets out an egress interface (for SAP egress queues). Defining a PIR does not necessarily guarantee that the queue can transmit at the intended rate. The actual rate sustained by the queue can be limited by oversubscription factors or available egress bandwidth.

The CIR defines the rate at which the system prioritizes the queue over other queues competing for the same bandwidth. In-profile, then out-of-profile packets are preferentially queued by the system at egress and at subsequent next hop nodes where the packet can traverse. To be properly handled throughout the network, the packets must be marked accordingly for profiling at each hop.

The CIR can be used by the queue's parent commands *cir-level* and *cir-weight* parameters to define the amount of bandwidth considered to be committed for the child queue during bandwidth allocation by the parent scheduler.

The **rate** command can be executed at any time, altering the PIR and CIR rates for all queues created through the association of the SAP egress QoS policy with the *queue-id*.

The **no** form of this command returns all queues created with the *queue-id* by association with the QoS policy to the default PIR and CIR parameters (**max**, 0).

Default rate max cir 0

Parameters *pir-rate* — Defines the administrative PIR rate, in kilobits, for the queue. When the **rate** command is executed, a valid PIR setting must be explicitly defined. When the **rate** command has not been executed, the default PIR of **max** is assumed.

Fractional values are not allowed and must be given as a positive integer.

The actual PIR rate is dependent on the queue's **adaptation-rule** parameters and the actual hardware where the queue is provisioned.

Values 1 to 100000000

Default max

cir-rate — The **cir** parameter overrides the default administrative CIR used by the queue. When the **rate** command is executed, a CIR setting is optional. When the **rate** command has not been executed or the **cir** parameter is not explicitly specified, the default CIR (0) is assumed.
Fractional values are not allowed and must be given as a positive integer. The **sum** keyword specifies that the CIR be used as the summed CIR values of the children schedulers or queues.

Values 0 to 100000000, **max**, **sum**

Default 0

slope-policy

Syntax **slope-policy** *hsmda-slope-policy-name*
no slope-policy

Context config>subscr-mgmt>sub-prof>hsmda>egress-qos>qos>queue

Description This command specifies an existing slope policy name. The policy contains the Maximum Buffer Size (MBS) that is applied to the queue and the high and low priority RED slope definitions. The function of the MBS and RED slopes is to provide congestion control for an HSMDA queue. The MBS parameter defines the maximum depth a queue may reach when accepting packets. The low and high priority RED slopes provides for random early detection of congestion and slope based discards based on queue depth.

An hsmda-slope-policy can be applied to queues defined in the sap-ingress and sap-egress QoS policy hsmda-queues context. Once an HSMDA slope policy is applied to a SAP QoS policy queue, it cannot be deleted. Any edits to the policy are updated to all HSMDA queues indirectly associated with the policy.

Default HSMDA Slope Policy

An hsmda-slope-policy named **default** always exists on the system and does not need to be created. The default policy is automatically applied to all HSMDA queues unless another HSMDA slope policy is specified for the queue. The default policy cannot be modified or deleted. Attempting to execute **no hsmda-slope-policy** default will result in an error.

The **no** form of this command removes the slope policy from the subscriber profile HSMDA configuration.

stat-mode

Syntax **stat-mode** {**v4-v6**}
no stat-mode

Context config>subscr-mgmt>sub-prof>hsmda>ingress-qos>qos>policer
config>subscr-mgmt>sub-prof>hsmda>ingress-qos>qos>queue

config>subscr-mgmt>sub-prof>hsmda>egress-qos>qos>queue

Description This command configures the forwarding plane octet and packet counters of a policer or queue to count packets of a specific type or state. For example separate counters for IPv4/IPv6.

For HSMDA ingress policers, this command overrides the policer stat-mode configuration as defined in the sap-ingress qos policy. For details on sap-ingress and sap-egress policer stat-mode, refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Quality of Service Guide*. For use in Enhanced Subscriber Management (ESM) context only, an additional stat-mode enables separate counters for IPv4 and IPv6 packets. **stat-mode v4-v6** is the only mode that can be configured as an HSMDA ingress policer override.

An HSMDA policer's stat-mode cannot be changed while the sub profile is in use.

For queues, this command sets the stat-mode. Queue stat-mode is only available for use in ESM context to enable separate IPv4/IPv6 counters.

An HSMDA queue's stat-mode cannot be changed while the sub profile is in use.

Default no stat-mode

For policers, the default is **no stat-mode override**. The sap-ingress stat-mode is used instead.

For queues, the default is to count in and out-of-profile octets and packets.

Parameters **v4-v6** — Count IPv4 and IPv6 forwarded/dropped octets and packets separately.

wrr-weight

Syntax **wrr-weight value**
no wrr-weight

Context config>subscr-mgmt>sub-prof>hsmda>egress-qos>qos>queue

Description This command assigns the weight value to the HSMDA queue.

The **no** form of this command returns the weight value for the queue to the default value.

Parameters *percentage* — Specifies the weight for the HSMDA queue.

Values 1 to 32

wrr-policy

Syntax **wrr-policy** *hsmda-wrr-policy-name*
no wrr-policy

Context config>subscr-mgmt>sub-prof>hsmda>egress-qos>qos

Description This command associates an existing HSMDA weighted-round-robin (WRR) scheduling loop policy to the HSMDA queue.

Parameters *hsmda-wrr-policy-name* — Specifies the existing HSMDA WRR policy name to associate to the queue.

9.25.2.1.25 Subscriber Management Service Commands

SAP Subscriber Management Commands

sub-sla-mgmt

Syntax [**no**] **sub-sla-mgmt**

Context config>service>vpls>sap
config>service>ies>if>sap
config>service>ies>sub-if>grp-if>sap
config>service>ies>sub-if>grp-if>sap-parameters
config>service>vprn>if>sap
config>service>vprn>sub-if>grp-if>sap
config>service>ies>sub-if>grp-if>sap-parameters

Description This command enables the context to configure subscriber management parameters for this SAP.

The **no** form of this command reverts to the default.

def-sla-profile

Syntax **def-sla-profile** *default-sla-profile-name*
no def-sla-profile

Context config>service>vpls>sap>sub-sla-mgmt
config>service>ies>if>sap>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap-param>sub-sla-mgmt
config>service>vprn>sub-if>grp-if>sap-param>sub-sla-mgmt

Description This command specifies a default SLA profile for this SAP.

An SLA profile is a named group of QoS parameters used to define per service QoS for all subscriber hosts common to the same subscriber within a provider service offering. A single SLA profile may define the QoS parameters for multiple subscriber hosts.

The **no** form of this command removes the default SLA profile from the SAP configuration.

Parameters *default-sla-profile-name* — Specifies a default SLA profile for this SAP.

def-sub-profile

Syntax **def-sub-profile** *default-subscriber-profile-name*

Context config>service>vpls>sap>sub-sla-mgmt
config>service>ies>if>sap>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap-param>sub-sla-mgmt
config>service>vprn>sub-if>grp-if>sap-param>sub-sla-mgmt

Description This command specifies a default subscriber profile for this SAP.

A subscriber profile defines the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for subscriber using the subscriber profile.

The **no** form of this command removes the default SLA profile from the SAP configuration.

Parameters *default-sub-profile* — Specifies a default subscriber profile for this SAP.

sub-ident-policy

Syntax **sub-ident-policy** *sub-ident-policy-name*

Context config>service>vpls>sap>sub-sla-mgmt
config>service>ies>if>sap>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap-param>sub-sla-mgmt
config>service>vprn>sub-if>grp-if>sap-param>sub-sla-mgmt

Description This command associates a subscriber identification policy to this SAP.

Subscribers are managed by the system through the use of subscriber identification strings. A subscriber identification string uniquely identifies a subscriber. For static hosts, the subscriber identification string is explicitly defined with each static subscriber host.

For dynamic hosts, the subscriber identification string must be derived from the DHCP ACK message sent to the subscriber host. The default value for the string is the content of Option 82 CIRCUIT-ID and REMOTE-ID fields interpreted as an octet sting. As an option, the DHCP ACK message may be processed by a subscriber identification policy which has the capability to parse the message into an alternative ASCII or octet string value.

When multiple hosts on the same port are associated with the same subscriber identification string they are considered to be host members of the same subscriber.

The **no** form of this command removes the default subscriber identification policy from the SAP configuration.

Parameters *sub-ident-policy-name* — Specifies a subscriber identification policy for this SAP,

multi-sub-sap

Syntax **multi-sub-sap** *number-of-sub*
no multi-sub-sap

Context config>service>vpls>sap>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt

Description This command defines the maximum number of subscribers (dynamic + static) that can be simultaneously active on this SAP.

If the limit is reached, a new host is denied access and the corresponding DHCP ACK is dropped.

The **no** form of this command reverts back to the default setting.

Parameters *multi-sub-sap* — Specifies the maximum number of subscribers allowed for this SAP. The operational maximum value may be smaller than the configured value due to equipped hardware dependencies.

Values 2 to 131071

single-sub-parameters

Syntax **single-sub-parameters**

Context config>service>ies>if>sap>sub-sla-mgmt>single-sub
config>service>ies>if>sap>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
config>service>vpls>sap>sub-sla-mgmt

Description This command enables the context to configure single subscriber SAP parameters.

non-sub-traffic

Syntax **non-sub-traffic sub-profile** *sub-profile-name* **sla-profile** *sla-profile-name* [**subscriber** *sub-ident-string*] [**app-profile** *app-profile-name*]

no non-sub-traffic

Context config>service>ies>if>sap>sub-sla-mgmt>single-sub
config>service>vpls>sap>sub-sla-mgmt>single-sub
config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt>single-sub

Description This command configures traffic profiles for non-IP traffic such as PPPoE. It is used in conjunction with the profiled-traffic-only on single subscriber SAPs and creates a subscriber host which is used to forward non-IP traffic through the single subscriber SAP without the need for SAP queues.

The **no** form of this command removes any configured profile.

Parameters *sub-profile-name* — Specifies an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

sla-profile-name — Specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

sub-ident-string — Specifies an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the **config>subscr-mgmt>sub-ident-policy** context. The subscriber information is used by the VPRN SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.

For SAPs with **arp-reply-agent** enabled with the optional *sub-ident* parameter, the static subscriber host's **sub-ident-string** is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the destinations.

If the static subscriber host's *sub-ident* string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.

If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. (ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.)

If *sub-ident* is not enabled on the SAP **arp-reply-agent**, subscriber identification matching is not performed on ARP requests received on the SAP.

ARP requests are never forwarded back to the same SAP or within the receiving SAP's Split Horizon Group.

app-profile-name — Specifies an existing app profile name.

profiled-traffic-only

Syntax [**no**] **profiled-traffic-only**

Context config>service>ies>if>sap>sub-sla-mgmt>single-sub
config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt>single-sub
config>service>vpls>sap>sub-sla-mgmt
config>service>vpls>sap>sub-sla-mgmt>single-sub-parameters
config>service>vprn>if>sap>sub-sla-mgmt>single-sub
config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt>single-sub

Description This command specifies whether only profiled traffic is applicable for this SAP. The profiled traffic refers to single subscriber traffic on a dedicated SAP (in the VLAN-per-subscriber model). When enabled, subscriber queues are instantiated through the QOS policy defined in the sla-profile and the associated SAP queues are deleted. This can increase subscriber scaling by reducing the number of queues instantiated per subscriber (in the VLAN-per-subscriber model). In order for this to be achieved, any configured multi-sub-sap limit must be removed (leaving the default of 1).

The **no** form of this command disables the command.

srrp

Syntax **srrp** *srrp-id* [**create**]
no srrp *srrp-id*

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command creates an SRRP instance on a group IP interface. An SRRP instance manages all subscriber subnets within the group interfaces subscriber IP interface or other subscriber IP interfaces that are associated through a wholesale/retail relationship. Only one unique SRRP instance can be configured per group interface.

The **no** form of this command removes an SRRP instance from a group IP interface. Once removed, the group interface ignores ARP requests for the SRRP gateway IP addresses that may exist on subscriber subnets associated with the group IP interface. Then the group interface stops routing using the redundant IP interface associated with the group IP interface and will stop routing with the SRRP gateway MAC address. Ingress packets destined to the SRRP gateway MAC will also be silently discarded. This is the same behavior as a group IP interface that is disabled (shutdown).

This command applies only to the 7750 SR.

The **no** form of this command removes the SRRP ID from the configuration.

Parameters *srrp-id* — Specifies a 32 bit instance ID that must be unique to the system. The instance ID must also match the instance ID used by the remote router that is participating in the same SRRP context. SRRP is intended to perform a function similar to VRRP where adjacent IP hosts within local subnets use a default gateway to access IP hosts on other subnets.

Values 1 to 4294967295

bfd-enable

Syntax [**no**] **bfd-enable** [*service-id*] **interface** *interface-name* **dst-ip** *ip-address*

Context config>service>ies>sub-if>grp-if>srrp
config>service>vprn>sub-if>grp-if>srrp

Description This commands assigns a bi-directional forwarding (BFD) session providing heart-beat mechanism for the given VRRP/SRRP instance. There can be only one BFD session assigned to any given VRRP/SRRP instance, but there can be multiple SRRP/VRRP sessions using the same BFD session.

BFD control the state of the associated interface. By enabling BFD on a given protocol interface, the state of the protocol interface is tied to the state of the BFD session between the local node and the remote node. The parameters used for the BFD are set via the BFD command under the IP interface.

The **no** form of this command disables BFD.

Parameters *service-id* — Specifies the service ID of the interface running BFD.

Values *service-id*: 1 to 214748364
svc-name: A string up to 64 characters

interface-name — Specifies the name of the interface running BFD.

ip-address — Specifies the destination address of the interface running BFD.

gw-mac

Syntax **gw-mac** *mac-address*
no gw-mac

Context config>service>vprn>sub-if>grp-if>srrp

Description This command overrides the default SRRP gateway MAC address used by the SRRP instance. Unless specified, the system uses the same base MAC address for all SRRP instances with the last octet overridden by the lower 8 bits of the SRRP instance ID. The same SRRP gateway MAC address should be in-use by both the local and remote routers participating in the same SRRP context.

One reason to change the default SRRP gateway MAC address is if two SRRP instances sharing the same broadcast domain are using the same SRRP gateway MAC. The system will use the SRRP instance ID to separate the SRRP messages (by ignoring the messages that does not match the local instance ID), but a unique SRRP gateway MAC is essential to separate the routed packets for each gateway IP address.

The **no** form of this command removes the explicit SRRP gateway MAC address from the SRRP instance. The SRRP gateway MAC address can only be changed or removed when the SRRP instance is shutdown.

Parameters *mac-address* — Specifies a MAC address that is used to override the default SRRP base MAC address.

Values Any MAC address except all zeros, broadcast or multicast addresses. The offset is expressed in normal Ethernet MAC address notation. The defined gw-mac cannot be 00:00:00:00:00:00, ff:ff:ff:ff:ff:ff or any multicast address.

If not specified, the system uses the default SRRP gateway MAC address with the last octet set to the 8 least significant bits of the SRRP instance ID.

keep-alive-interval

Syntax **keep-alive-interval** *interval*
no keep-alive-interval

Context config>service>ies>sub-if>grp-if>srrp
config>service>vprn>sub-if>grp-if>srrp

Description This command defines the interval between SRRP advertisement messages sent when operating in the master state. The interval is also the basis for setting the master-down timer used to determine when the master is no longer sending. The system uses three times the keep-alive interval to set the timer. Every time an SRRP advertisement is seen that is better than the local priority, the timer is reset. If the timer expires, the SRRP instance assumes that a master does not exist and initiates the attempt to become master.

When in backup state, the SRRP instance takes the keep-alive interval of the master as represented in the masters SRRP advertisement message. Once in master state, the SRRP instance uses its own configured keep-alive interval.

The keep-alive-interval may be changed at anytime, but will have no effect until the SRRP instance is in the master state.

The **no** form of this command restores the default interval.

Default keep-alive-interval 10

Parameters *interval* — Specifies the interval, in deciseconds, between SRRP advertisement messages sent when operating in the master state.

Values 1 to 100

message-path

Syntax **message-path** *sap-id*
no message-path

Context config>service>ies>sub-if>grp-if>srrp
config>service>vprn>sub-if>grp-if>srrp

Description This command defines a specific SAP for SRRP in-band messaging. A message-path SAP must be defined prior to activating the SRRP instance. The defined SAP must exist on the SRRP instances group IP interface for the command to succeed and cannot currently be associated with any dynamic or static subscriber hosts. Once a group IP interface SAP has been defined as the transmission path for SRRP Advertisement messages, it cannot be administratively shutdown, will not support static or dynamic subscriber hosts and cannot be removed from the group IP interface.

The SRRP instance message-path command may be executed at anytime on the SRRP instance. Changing the message SAP will fail if a dynamic or static subscriber host is associated with the new SAP. Once successfully changed, the SRRP instance will immediately disable anti-spoof on the SAP and start sending SRRP Advertisement messages if the SRRP instance is activated.

Changing the current SRRP message SAP on an active pair of routers should be done in the following manner:

1. Shutdown the backup SRRP instance.
2. Change the message SAP on the shutdown node.
3. Change the message SAP on the active master node.
4. Re-activate the shutdown SRRP instance.

Shutting down the backup SRRP instance prevents the SRRP instances from becoming master due to temporarily using differing message path SAPs.

If an MCS peering is operational between the redundant nodes and the SRRP instance has been associated with the peering, the designated message path SAP is sent from each member.

The **no** form of this command can only be executed when the SRRP instance is shutdown. Executing no message-path allows the existing SAP to be used for subscriber management functions. A new message-path SAP must be defined prior to activating the SRRP instance.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

monitor-oper-group

Syntax **monitor-oper-group** *name* **priority-step** *step*
no monitor-oper-group

Context config>service>ies>sub-if>grp-if>srrp
config>service>vprn>sub-if>grp-if>srrp

Description This command will configure the association between the SRRP instance in a Fate Sharing Group and the operational-group that contains messaging SAPs. A state transition of a messaging SAP within an operational-group will trigger calculation of the priority for all SRRP instances that are associated with that operational-group.

The **no** form of this command reverts to the default.

Parameters *name* — Specifies the name of the operational-group, up to 32 characters, that is tracking operational state of SRRP messaging SAPs

priority-step — Specifies the priority step for which the priority of an SRRP instance is changed. If a messaging SAP within an operational-group transition to a non-UP state, the priority is decreased by the step value. If the messaging SAP within the operational-group transition into the UP state, the priority of the SRRP instance is increased by the step value.

Values 1 to 10

one-garp-per-sap

Syntax [**no**] **one-garp-per-sap**

Context config>service>ies>sub-if>grp-if>srrp
config>service>vprn>sub-if>grp-if>srrp

Description This command enables the sending of one gratuitous ARP to each SAP and is applicable to PPPoE only deployments in which there are multiple subnets under the subscriber interface. In such case, if the switchover occurs, it is sufficient to send a single Gratuitous ARP on every VLAN to update the Layer 2 forwarding path in the access aggregation network. This single gratuitous ARP will contain the IP address of the first GW address found under the subscriber interface address.

The **no** form of this command reverts to the default.

policy

Syntax [**no**] **policy** *vrrp-policy-id*

Context config>service>ies>sub-if>grp-if>srrp
config>service>vprn>sub-if>grp-if>srrp

Description This command associates one or more VRRP policies with the SRRP instance. A VRRP policy is a collection of connectivity and verification tests used to manipulate the in-use priorities of VRRP and SRRP instances. A VRRP policy can test the link state of ports, ping IP hosts, discover the existence of routes in the routing table or the ability to reach Layer 2 hosts. When one or more of these tests fail, the VRRP policy has the option of decrementing or setting an explicit value for the in-use priority of an SRRP instance.

More than one VRRP policy may be associated with an SRRP instance. When more than one VRRP policy is associated with an SRRP instance the delta decrement of the in-use priority is cumulative unless one or more test fail that have explicit priority values. When one or more explicit tests fail, the lowest priority value event takes effect for the SRRP instance. When the highest delta-in-use-limit is used to manage the lowest delta derived in-use priority for the SRRP instance.

VRRP policy associations may be added and removed at anytime. A maximum of two VRRP policies can be associated with a single SRRP instance.

The **no** form of this command removes the association with vrrp-policy-id from the SRRP instance.

Parameters *vrrp-policy-id* — Specifies one or more VRRP policies with the SRRP instance.

Values 1 to 9999

preempt

Syntax [**no**] **prempt**

Context config>service>ies>sub-if>grp-if>srrp
config>service>vprn>sub-if>grp-if>srrp

Description When preempt is enabled, a newly initiated SRRP instance can overrides an existing Master SRRP instance if its priority value is higher than the priority of the current Master.

If preempt is disabled, an SRRP instance only becomes Master if the master down timer expires before an SRRP advertisement message is received from the adjacent SRRP enabled node.

The **no** form of this command reverts to the default.

Default preempt

priority

Syntax **priority** *priority*
no priority

Context config>service>ies>sub-if>grp-if>srrp

config>service>vprn>sub-if>grp-if>srrp

Description This command overrides the default base priority for the SRRP instance. The SRRP instance priority is advertised by the SRRP instance to its neighbor router and is compared to the priority received from the neighbor router. The router with the best (highest) priority enters the master state while the other router enters the backup state. If the priority of each router is the same, the router with the lowest source IP address in the SRRP advertisement message assumes the master state.

The base priority of an SRRP instance can be managed by VRRP policies. A VRRP policy defines a set of connectivity or verification tests which, when they fail, may lower an SRRP instances base priority (creating an in-use priority for the instance). Every time an SRRP instances in-use priority changes when in master state, it sends an SRRP advertisement message with the new priority. If the dynamic priority drops to zero or receives an SRRP Advertisement message with a better priority, the SRRP instance transitions to the *becoming backup* state.

When the priority command is not specified, or the no priority command is executed, the system uses a default base priority of 100. The priority command may be executed at anytime.

The **no** form of this command restores the default base priority to the SRRP instance. If a VRRP policy is associated with the SRRP instance, it will use the default base priority as the basis for any modifications to the SRRP instances in-use priority.

Default priority 100

Parameters *priority* — Specifies a base priority for the SRRP instance to override the default.

Values 1 to 254

send-fib-population-packets

Syntax **send-fib-population-packets** *mode*
no send-fib-population-packets

Context config>service>ies>sub-if>grp-if>srrp
config>service>vprn>sub-if>grp-if>srrp

Description This command configures the mode used to send Fib population packets. When SRRP becomes master it generates gratuitous ARPs (GARPs) used by the Layer 2 access network to populate the correct SRRP gateway.

The **no** form of this command disables sending FDB population packets.

Default send-fib-population-packets all

Parameters *mode* — Specifies on which VLANs the gratuitous ARPs are sent.

Values **all**: Generates, on SSRP master assignment, the GARPs on all VLANs

out-tag-only: Generates, on SRRP master assignment, the GARPs only on SAPs with unique outer VLAN and lowest VLAN tags

srrp-enabled-routing

Syntax **srrp-enabled-routing** [**hold-time** *hold-time*]
no srrp-enabled-routing

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command configures SRRP-enabled routing.

The **no** form of this command reverts to the default.

Parameters **hold-time** *hold-time* — Specifies the hold time, in deci-seconds.

Values 1 to 50

tos-marking-state

Syntax **tos-marking-state** {**trusted** | **untrusted**}
no tos-marking-state

Context config>service>ies>if
config>service>ies>sub-if>grp-if
config>service>vprn>interface
config>service>vprn>sub-if>grp-if

Description This command is used to alter the default trusted state to a non-trusted state. When unset or reverted to the trusted default, the ToS field will not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set, in which case the egress network interface treats all VPRN and network IP interface as untrusted.

When the ingress interface is set to untrusted, all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface. The egress network remarking rules also apply to the ToS field of IP packets routed using IGP shortcuts (tunneled to a remote next-hop). However, the tunnel QoS markings are always derived from the egress network QoS definitions. Egress marking and remarking is based on the internal forwarding class and profile state of the packet once it reaches the egress interface. The forwarding class is derived from ingress classification functions. The profile of a packet is either derived from ingress classification or ingress policing.

The default marking state for network IP interfaces is trusted. This is equivalent to declaring no tos-marking-state on the network IP interface. When undefined or set to tos-marking-state trusted, the trusted state of the interface will not be displayed when using show config or show info unless the detail parameter is given. The **save config** command will not store the default tos-marking-state trusted state for network IP interfaces unless the detail parameter is also specified.

The **no** form of this command restores the trusted state to a network IP interface. This is equivalent to executing the tos-marking-state trusted command.

Default tos-marking-state trusted

Parameters **trusted** — The default prevents the ToS field to not be remarked by egress network IP interfaces unless the egress network IP interface has the remark-trusted state set.

untrusted — Specifies that all egress network IP interfaces will remark IP packets received on the network interface according to the egress marking definitions on each network interface.

mac-da-hashing

Syntax **mac-da-hashing**
no mac-da-hashing

Context config>service>vpls>sap>sub-sla-mgmt

Description This command specifies whether subscriber traffic egressing a LAG SAP has its egress LAG link selected by a function of the MAC destination address instead of the subscriber ID.

The **no** form of this command reverts to the default setting.

diameter-auth-policy

Syntax **diameter-auth-policy** *name*
no diameter-auth-policy

Context config>service>vpls>sap

Description This command is used to configure the Diameter NASREQ application policy to use for authentication.

The **no** form of this command reverts to the default.

Parameters *name* — Specifies the name of the Diameter NASREQ application policy to use for authentication.

host

Syntax **host** {[**ip** *ip-address* [**mac** *mac-address*]} [**subscriber-sap-id** | **subscriber** *sub-ident-string* [**sub-profile** *sub-profile-name* [**sla-profile** *sla-profile-name* [**ancp-string** *ancp-string*] [**app-profile** *app-profile-name*] [**inter-dest-id** *intermediate-destination-id*]

no host {[**ip** *ip-address*] [**mac** *ieee-address*]}

no host all

Context config>service>vpls>sap
config>service>ies>sub-if>grp-if>sap
config>service>ies>if>sap
config>service>vprn>sub-if>grp-if>sap

Description This command creates a static subscriber host for the SAP. Static subscriber hosts may be used by the system for various purposes. Applications within the system that make use of static host entries include anti-spoof, ARP reply agent and source MAC population into the VPLS forwarding database.

Multiple static hosts may be defined on the SAP. Each host is identified by either a source IP address, a source MAC address or both a source IP and source MAC address. Every static host definition must have at least one address defined, IP or MAC.

Static hosts can exist on the SAP even with anti-spoof and ARP reply agent features disabled. When enabled, each feature has different requirements for static hosts.

The **no** form of this command removes a static entry from the system. The specified *ip-address* and *mac-address* must match the host's exact IP and MAC addresses as defined when it was created. When a static host is removed from the SAP, the corresponding anti-spoof filter entry and/or FDB entry is also removed.

Parameters *ip-address* — Specifies an association with a subscriber with the static subscriber host. Only one static host can be configured on the SAP with a given IP address.

mac-address — Specifies the MAC address to define a static host. The MAC address must be specified for **anti-spoof mac**, **anti-spoof ip-mac** and **arp-reply-agent** (the **arp-reply-agent** is supported by the 7450 ESS only). Multiple static hosts may be configured with the same MAC address given that each definition is distinguished by a unique IP address.

Every static host definition must have at least one address defined, IP or MAC.

sub-ident-string — Specifies an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the **config>subscr-mgmt>sub-ident-pol** context. The subscriber information is used by the VPLS SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.

For VPLS SAPs with **arp-reply-agent** enabled with the optional *sub-ident* parameter, the static subscriber hosts sub-ident-string is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the VPLS destinations.

If the static subscriber hosts *sub-ident* string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.

If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. (ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.)

If *sub-ident* is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.

ARP requests are never forwarded back to the same SAP or within the receiving SAP's Split Horizon Group.

sub-profile *sub-profile-name* — Specifies an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

sla-profile *sla-profile-name* — Specifies an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

Note: If Enhanced Subscriber Management is enabled on a SAP using the **sub-sla-mgmt** command, the **sub-ident**, **sub-profile,** and **sla-profile** must be configured for all static hosts defined on this SAP.

Wireless Portal Protocol (WPP) Commands

wpp

Syntax **wpp**

Context config>aaa
config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command enables the context to configure Wireless Portal Protocol (WPP) parameters.

portal-group

Syntax **portal-group** *portal-group-name* [**create**]

no portal-group *portal-group-name*

Context config>aaa>wpp>portal-groups

Description This command creates a new portal group or enters the configuration context of an existing port group.

Parameters *portal-group-name* — Specifies the portal group name up to 32 characters

create — Keyword required to create the configuration context.

portal-group

Syntax **portal-group** *portal-group-name*
no portal-group

Context config>service>ies>sub-if>grp-if>wpp
config>service>vprn>sub-if>grp-if>wpp

Description This command specifies the WPP portal group for the subscriber interface. This command is mutually exclusive with the **portal** command.

The **no** form of this command removes the name from the service configuration.

Parameters *portal-group-name* — Specifies the portal group name up to 32 characters.

portal-group

Syntax **portal-group** *portal-group-name*
no portal-group

Context config>subscr-mgmt>loc-user-db>ipoe>host>wpp

Description This command specifies the WPP portal group for the LUDB host. This command is mutually exclusive with the **portal** command.

The **no** form of this command removes the name from the LUDB configuration.

Parameters *portal-group-name* — Specifies the portal group name up to 32 characters.

portal-groups

Syntax **portal-groups**

Context config>aaa>wpp

Description This command enables the context to configure portal group parameters for WPP.

portal

Syntax [**no**] **portal router** *router-instance* **name** *wpp-portal-name*

Context config>aaa>wpp>portal-groups>portal-group

Description This command configures the portal for this portal group.

Parameters *router-instance* — Specifies the virtual router instance.

Values

router-name:	Base, management
service-id:	1 to 2147483647
service-name:	Specifies the service name up to 64 characters

Default Base

wpp-portal-name — Specifies the name of the web portal server up to 32 characters.

portal

Syntax **portal router** *router-instance* **name** *wpp-portal-name*
no portal

Context config>subscr-mgmt>loc-user-db>ipoe>host>wpp
config>service>ies>sub-if>grp-if>wpp
config>service>vprn>sub-if>grp-if>wpp

Description This command specifies the web portal server that system talks to for the hosts on the group-interface. This command is mutually exclusive with the **portal-group** command.

The **no** form of this command removes the router instance or portal name from the configuration.

Parameters *router-instance* — Specifies the virtual router instance.

Values

router-name:	Base, management
service-id:	1 to 2147483647
service-name:	Specifies the service name up to 64 characters

Default Base

wpp-portal-name — Specifies the name of the web portal server up to 32 characters.

portals

Syntax **portals**

Context config>router>wpp
config>service>vprn>wpp

Description This command enables the context to configure WPP portal server parameters.

ack-auth-retry-count

Syntax **ack-auth-retry-count** [*value*]
no ack-auth-retry-count

Context config>router>wpp>portals>portal
config>service>vprn>wpp>portals>portal

Description This command configures the number of retransmissions of an ACK_OUT message.

The **no** form of this command reverts to the default.

Default ack-auth-retry-count 5

Parameters *value* — Specifies the number of retransmissions of an ACK_OUT message.

Values 0 to 5

ntf-logout-retry-count

Syntax **ntf-logout-retry-count** [*value*]
no ntf-logout-retry-count

Context config>router>wpp>portals>portal
config>service>vprn>wpp>portals>portal

Description This command configures the number of retransmissions of an NTF_LOGOUT message.

The **no** form of this command reverts to the default.

Default ntf-logout-retry-count 5

Parameters *value* — Specifies the number of retransmissions of an NTF_LOGOUT message.

Values 0 to 5

port-format

Syntax **port-format** *formatting*
no port-format

Context config>router>wpp>portals>portal
config>service>vprn>wpp>portals>portal

Description This command specifies the encoding format of WPP port attribute.

The standard format is as follows:

<0 to 20 character system-name><1 character separator><2-digit slot><1-digit mda><2-digit port><4-digit top><5-digit bottom>

As a general rule, if a value is not present or is too large to fit in the field, is the field set to all zeros. The following rules apply to standard formats.

- With a standard port, when the separator is a “-” character, the slot is the *slot-id*, mda is the *mda-id*, and the port is the *port-id*.
- With an ESAT port, when the separator sep is a “:” character, the slot is the *satellite-id*, MDA is satellite *slot-id*, and the port is satellite *port-id*.
- With a PXC port, when separator is a “#” character, the MDA is the PXC *subport-id*, and the port is the PXC *port-id*.
- With a LAG port, the port is the *lag-id*.
- With a connector port, the slot is the *slot-id*, the MDA is the *mda-id*, and the port is the *connector-id*.

The vendor-specific format is as follows:

With dot1q, append “%u” with the top *vlan-id*.

With qinq, append “%u.%u” with the top *vlan-id* and the bottom *vlan-id*.

As a general rule, there can be no trailing zeros. The string truncates if it becomes too long. 0 to 16 characters are allowed for the system name. The following rules apply to vendor-specific formats.

- With a standard port, append “%s-%u/%u/%” with the *system-name*, *slot-id*, *mda-id*, and *port-id*.
- With an ESAT port, append “%s-S%u/%u/%u” with the *system-name*, *satellite-id*, *satellite-slot-id*, and *satellite-port-id*.
- With a PXC port, append “%s-P%u%c” with the *system-name*, PXC *port-id*, and PXC *subport-id* ? 'a' : 'b'.
- With a LAG port, append “%s-L%u” with the *system-name* and *lag-id*.
- With a connector port append “%s-%u%uc%u/%u” with the *system-name*, *slot-id*, *mda-id*, *connector-id*, and *connector-port-id*.

The **no** form of this command reverts to the default.

Default port-format standard

Parameters *formatting* — Specifies the encoding format of the WPP port attribute.

Values standard, vendor-specific

retry-interval

Syntax **retry-interval** *milliseconds*
no retry-interval

Context config>router>wpp>portals>portal
config>service>vprn>wpp>portals>portal

Description This command configures the time interval between two consecutive retransmissions

The **no** form of this command reverts to the default.

Default retry-interval 2000

Parameters *milliseconds* — Specifies the time interval between two consecutive retransmissions.

Values 10 to 2000

secret

Syntax **secret** *secret* [**hash** | **hash2**]
no secret

Context config>router>wpp>portals>portal
config>service>vprn>wpp>portals>portal

Description This command configures the secret that is used by WPPv2 to authenticate the messages between portal and BRAS.

The **no** form of this command removes the secret and hash from the configuration.

Parameters *secret* — Specifies the secret key to access the server, up to 64 characters.

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

version

Syntax **version** *version*
no version

Context config>router>wpp>portals>portal
config>service>vprn>wpp>portals>portal

Description This command configure the protocol version that is expected by the WPP portal.

The **no** form of this command reverts to the default.

Default version 1

Parameters *version* — Specifies the protocol version.

Values 1, 2

enable-triggered-hosts

Syntax [**no**] **enable-triggered-hosts**

Context config>service>vprn>sub-if>grp-if>wpp
config>service>ies>sub-if>grp-if>wpp

Description This command enables system to auto creates ESM hosts upon successful WPP authentication. The default host must be configured under SAP on the subscriber SAP to redirect unauthenticated client traffic to the web portal.

The **no** form of this command reverts to the default.

initial-app-profile

Syntax **initial-app-profile** *app-profile-name*
no initial-app-profile

Context config>subscr-mgmt>loc-user-db>ipoe>host>wpp
config>service>ies>sub-if>grp-if>wpp
config>service>vprn>sub-if>grp-if>wpp

Description This command specifies the initial app-profile for the hosts created on the group-interface. This initial app-profile is replaced after hosts pass web portal authentication.

The **no** form of this command reverts to the default.

Parameters *app-profile-name* — Specifies the initial application profile, up to 32 characters, to be used during the WPP authentication phase of the IPoE hosts.

initial-sla-profile

Syntax **initial-sla-profile** *sla-profile-name*
no initial-sla-profile

Context config>subscr-mgmt>loc-user-db>ipoe>host>wpp
config>service>ies>sub-if>grp-if>wpp
config>service>vprn>sub-if>grp-if>wpp

Description This command specifies the initial sla-profile for the hosts created on the group-interface. This initial sla-profile is replaced after hosts pass web portal authentication.

The **no** form of this command reverts to the default.

Parameters *sla-profile-name* — Specifies the initial SLA profile, up to 32 characters, to be used during the WPP authentication phase of the IPOE host.

initial-sub-profile

Syntax **initial-sub-profile** *sub-profile-name*
no initial-sub-profile

Context config>subscr-mgmt>loc-user-db>ipoe>host>wpp
config>service>ies>sub-if>grp-if>wpp
config>service>vprn>sub-if>grp-if>wpp

Description This command specifies the initial sub-profile for the hosts created on the group-interface. This initial sub-profile is replaced after hosts pass web portal authentication.

The **no** form of this command reverts to the default.

Parameters *sub-profile-name* — Specifies the initial subscriber profile, up to 32 characters, to be used during the WPP authentication phase of the IPoE host.

lease-time

Syntax **lease-time** [*lease-time*]
no lease-time

Context config>service>ies>sub-if>grp-if>wpp

Description This command configures the amount of time that the DHCP server grants to the DHCP client permission to use a particular IP address.

The **no** form of this command removes the lease time parameters from the configuration.

Parameters *lease-time* — Specifies the lease time.

Values

days *days*	0 to 3650
hrs *hours*	0 to 23
min *minutes*	0 to 59
sec *seconds*	0 to 59

restore-disconnected

Syntax [**no**] **restore-disconnected**

Context config>service>ies>sub-if>grp-if>wpp
config>service>vprn>sub-if>grp-if>wpp

Description This command specifies that the initial profiles must be restored after a DHCP host has disconnected. The behavior that system will restore the **initial-sla-profile**, **initial-sub-profile**, or **initial-app-profile** when hosts disconnects instead of removing them.

The **no** form of this command specifies that the initial profiles will not be restored after a DHCP host has disconnected.

Default restore-disconnected

restore-disconnected

Syntax **restore-disconnected** {**restore** | **no-restore**}
no restore-disconnected

Context config>subscr-mgmt>loc-user-db>ipoe>host>wpp

Description This command specifies the behavior that system will restore the **initial-sla-profile**, **initial-sub-profile**, or **initial-app-profile** when hosts disconnects instead of removing them.

The **no** form of this command reverts to the default.

Parameters **restore** — Specifies that the initial profiles must be restored after a DHCP host has disconnected.

no-restore — Specifies that the initial profiles will not be restored after a DHCP host has disconnected.

user-db

Syntax **user-db** *local-user-db-name*
no user-db

Context config>subscr-mgmt>loc-user-db>ipoe>host>wpp
config>service>ies>sub-if>grp-if>wpp
config>service>vprn>sub-if>grp-if>wpp

Description This command configures the user database.

Note: If configured, the values configured under grp-if will only be used if there is no corresponding value returned from LUDB lookup.

This command specifies the LUDB system use to lookup while creating initial host before WPP authentication. LUDB could return WPP attributes such as portal name, **initial-sla-profile**, **initial-sub-profile**, and so on LUDB is configured in **config>subscr-mgmt>local-user-db** context.

The **no** form of this command reverts to the default.

Parameters *local-user-db-name* — Specifies the Local User Database name up to 32 characters.

9.25.2.1.26 Subscriber Management Service Commands

Note: The commands described in this section apply only to the 7750 SR.

subscriber-interface

Syntax **subscriber-interface** *ip-int-name* [**create**] **fwd-service** *service-id* **fwd-subscriber-interface** *ip-int-name* [**wan-mode** *mode*]
no subscriber-interface *ip-int-name*

Context config>service>ies
config>service>vprn

Description This command allows the operator to create special subscriber-based interfaces. It is used to contain multiple group interfaces. Multiple subnets associated with the subscriber interface can be applied to any of the contained group interfaces in any combination. The subscriber interface allows subnet sharing between group interfaces.

The **no** form of this command reverts to the default.

Parameters *ip-int-name* — Specifies the interface name of a subscriber interface. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

service-id — Specifies the wholesale service ID or service name.

Values *service-id*: 1 to 214748364
svc-name: A string up to 64 characters

ip-int-name — Specifies the wholesale subscriber interface.

wan-mode *mode* — Specifies the WAN mode as 64-bit or 128-bit. To change the WAN mode after creation, the interface must first be removed then recreated.

Values *mode64*, *mode128*

Default *mode64*

address

Syntax **address** *ipv6-address/prefix-length* [**eui-64**] [**track-srrp** *srrp-instance*] [**modifier** *cga-modifier*] [**dad-disable**]
no address *ipv6-address/prefix-length*

Context config>service>ies>if>ipv6
config>service>vprn>sub-if>ipv6

Description This command assigns an IPv6 address/subnet to the subscriber interface.

Note: SRRP is not supported for IPv6 subscriber interface.

The **no** form of this command reverts to the default.

Parameters *ipv6-address* — Specifies the 128-bit IPv6 address.

Values 128-bit hexadecimal IPv6 address in compressed form

prefix-length — Specifies the length of any associated aggregate prefix.

Values 32 to 63

srrp-id — Specifies the SRRP instance number.

Values 1 to 4294967295

cgi-modifier — Specifies the CGI modifier.

Values 0x0..0xFFFFFFFF...(32 hex nibbles)

dad-disable — Disables Duplicate Address Detection (DAD).

address

Syntax **address** *ipv6-address/prefix-length* [**pd**] [**wan-host**] [**track-srrp** *srrp-instance*] [**holdup-**

time *milli-seconds*]
no address *ipv6-address/prefix-length*

Context config>service>ies>sub-if>ipv6
config>service>vprn>sub-if>ipv6

Description This command assigns an IPv6 address/subnet to the subscriber interface.

The **no** form of this command reverts to the default.

Parameters *ipv6-address* — Specifies the 128-bit IPv6 address.

Values 128-bit hexadecimal IPv6 address in compressed form

prefix-length — Specifies the length of any associated aggregate prefix.

Values 32 to 63

pd — Specifies that this aggregate is used by IPv6 ESM hosts for DHCPv6 prefix-delegation.

wan-host — Specifies that this aggregate is used by IPv6 ESM hosts for local addressing or by a routing gateway's WAN interface.

srrp-instance — Specifies the SRRP instance number.

Values 1 to 4294967295

milli-seconds — Specifies the time to wait, in milli-seconds, for the route before it accepts the new state attribute. This timer is used to prevent fluctuations in route advertisement caused by short lived SRRP instabilities, in the case that such condition arises.

Values 100 to 5000

allow-unmatching-prefixes

Syntax [**no**] **allow-unmatching-prefixes**

Context config>service>ies>sub-if>ipv6
config>service>vprn>sub-if>ipv6

Description This command allows address assignment for IPoEv6 and PPPoEv6 hosts in cases where the subscriber host assigned IPv6 address or prefix falls outside of the subscriber-prefix range explicitly configured for the subscriber-interface (**configure>service>vprn/ies>sub-if>ipv6**) or the subscriber-prefix is not configured at all.

SLAAC hosts is installed in the FDB as /64 entries, the length of the installed DHCP-PD prefix is dictated by the prefix-length and the DHCP-NA host is installed as /128 entries.

IPv4 subscriber hosts are unaffected by this command.

The **no** form of this command reverts to the default.

arp-host

Syntax **arp-host**

Context config>service>ies>sub-if>grp-if
config>service>vpls>sap>arp-host
config>service>vprn>sub-if>grp-if
config>subscr-mgmt>msap-policy>vpls-only

Description This command enables the context to configure ARP host parameters.

authentication-policy

Syntax **authentication-policy** *name*
no authentication-policy

Context config>service>vprn>if
config>service>vprn>sub-if>grp-if

Description This command assigns an authentication policy to the interface.

The **no** form of this command removes the policy name from the group interface configuration.

Parameters *name* — Specifies the authentication policy name. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

host-limit

Syntax **host-limit** *max-num-hosts*
no host-limit

Context config>subscr-mgmt>msap-policy>vpls-only>arp-host
config>service>vpls>sap>arp-host
config>service>ies>sub-if>grp-if>arp-host
config>service>vprn>sub-if>grp-if>arp-host

Description This command configures the maximum number of ARP hosts.

The **no** form of this command reverts to the default.

Default host-limit 1

Parameters *max-num-hosts* — Specifies the maximum number of ARP hosts.

Note: The operational maximum value may be smaller due to equipped hardware dependencies.

Values 1 to 131071

link-local-address

Syntax **link-local-address** *ipv6-address*
no link-local-address

Context config>service>ies>if>ipv6
config>service>ies>sub-if>ipv6
config>service>vprn>if>ipv6
config>service>vprn>sub-if>ipv6

Description This command configures the IPv6 Link Local address that is used as a virtual SRRP IPv6 address by the Master SRRP node. This address is sent in the Router Advertisements initiated by the Master SRRP node. Clients use this address as IPv6 default-gateway. Both SRRP nodes, Master and Backup, must be configured with the same Link Local address.

The **no** form of this command reverts to the default.

Parameters *ipv6-address* — Specifies the IPv6 address in the form:

x:x:x:x:x:x:x:x:x:x:x:x:d.d.d.d
where:
x - [0 to FF]
d - [0 to 255]

prefix

Syntax **prefix** *ipv6-address/prefix-length* [**pd**] [**wan-host**]
no prefix *ipv6-address/prefix-length*

Context config>services>ies>sub-if>ipv6>sub-prefixes
config>services>vprn>sub-if>ipv6>sub-prefixes

Description This command allows a list of prefixes (using the prefix command multiple times) to be routed to hosts associated with this subscriber interface. Each prefix is represented in the associated FDB with a reference to the subscriber interface. Prefixes are defined as being for prefix delegation (pd) or use on a WAN interface or host (wan-host).

The **no** form of this command reverts to the default.

Parameters *ipv6-address* — Specifies the 128-bit IPv6 address.

Values 128-bit hexadecimal IPv6 address in compressed form

prefix-length — Specifies the length of any associated aggregate prefix.

Values 32 to 63

pd — Specifies that this aggregate is used by IPv6 ESM hosts for DHCPv6 prefix-delegation.

wan-host — Specifies that this aggregate is used by IPv6 ESM hosts for local addressing or by a routing gateway's WAN interface.

auto-reply

Syntax [**no**] **auto-reply**

Context config>service>vprn>sub-if>grp-if>ipv6
config>service>ies>sub-if>grp-if>ipv6

Description This command assists IP-only static hosts to resolve their default gateway and MAC. By default, the BNG anti-spoof filter drops packets from unknown hosts. The auto-reply features first allow hosts to resolve their default gateway and afterwards allow them to forward traffic. Using the data traffic, the BNG can utilize the data-trigger mechanism to learn the host's MAC and populate the full IP+MAC static host entry.

The **no** form of this command reverts to the default.

neighbor-solicitation

Syntax [**no**] **neighbor-solicitation**

Context config>service>vprn>sub-if>grp-if>ipv6>auto-reply
config>service>ies>sub-if>ipv6>auto-reply

Description This command enables auto-reply for neighbor solicitation.

The **no** form of this command disables auto-reply.

Default neighbor-solicitation

router-solicitation

Syntax [**no**] **router-solicitation**

Context config>service>vprn>sub-if>grp-if>ipv6>auto-reply
config>service>ies>sub-if>grp-if>ipv6>auto-reply

Description This command enables auto-reply router solicitation.

The **no** form of this command disables router-solicitation.

Default router-solicitation

arp-populate

Syntax [**no**] **arp-populate**

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command enables populating static and dynamic hosts into the system ARP cache. When enabled, the host's IP address and MAC address are placed in the system ARP cache as a managed entry. Static hosts must be defined on the interface using the **host** command. Dynamic hosts are enabled on the system through enabling lease-populate in the IP interface DHCP context. In the event that both a static host and a dynamic host share the same IP and MAC address, the system's ARP cache retains the host information until both the static and dynamic information are removed. Both static and dynamic hosts override static ARP entries. Static ARP entries are marked as inactive when they conflict with static or dynamic hosts and is repopulated once all static and dynamic host information for the IP address are removed. Since static ARP entries are not possible when static subscriber hosts are defined or when DHCP lease state table population is enabled, conflict between static ARP entries and the arp-populate function is not an issue.

The **arp-populate** command fails if an existing static subscriber host on the SAP does not have both MAC and IP addresses specified.

Once **arp-populate** is enabled, creating a static subscriber host on the SAP without both an IP address and MAC address fails.

arp-populate can only be enabled on VPRN interfaces supporting Ethernet encapsulation.

The **no** form of this command disables ARP cache population functions for static and dynamic hosts on the interface. All static and dynamic host information in the systems ARP cache is removed. Any existing static ARP entries previously inactive due to static or dynamic hosts are populated in the system ARP cache.

When **arp-populate** is enabled, the system will not send out ARP Requests for hosts that are not in the ARP cache. Only statically configured and DHCP learned hosts are reachable through an IP interface with arp-populate enabled.

arp-timeout

Syntax **arp-timeout** *seconds*
no arp-timeout

Context config>service>ies>if
config>service>ies>sub-if>grp-if
config>service>vprn>if
config>service>vprn>sub-if>grp-if

Description This command configures the minimum time in seconds an ARP entry learned on the IP interface is stored in the ARP table. ARP entries are automatically refreshed when an ARP request or gratuitous ARP is seen from an IP host, otherwise, the ARP entry is aged from the ARP table. If **arp-time-out** is set to a value of zero seconds, ARP aging is disabled.

The **no** form of this command restores **arp-timeout** to the default value.

Default arp-timeout 14400

Parameters *seconds* — Specifies the minimum number of seconds a learned ARP entry is stored in the ARP table, expressed as a decimal integer. A value of zero specifies that the timer is inoperative and learned ARP entries will not be aged.

Values 0 to 65535

authentication-policy

Syntax **authentication-policy** *name*
no authentication-policy

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command assigns a RADIUS authentication policy to the interface.

The **no** form of this command removes the policy name from the group interface configuration.

Parameters *name* — Specifies the authentication policy name.

export-host-routes

Syntax [**no**] **export-host-routes**

Context config>service>ies>subscriber-interface
config>service>vprn>subscriber-interface

Description This command controls the export of subscriber management host routes from a retail service to the corresponding forwarding wholesale VPRN service.

By default, subscriber management host routes are not exported.

The presence of retail subscriber management host routes in the wholesale VPRN service is required for downstream traffic forwarding in multi-chassis redundancy scenario's with a redundant interface and when the retail subscriber subnets are not leaked in the wholesale VPRN service (allow-unmatching-subnets or unnumbered retail subscriber interface).

This command fails if the subscriber interface is not associated with a forwarding wholesale service subscriber interface or if the subscriber interface is not configured to support address allocation outside the provisioned subnets (allow-unmatching-subnets or unnumbered subscriber interface)

The **no** form of this command reverts to the default.

group-interface

Syntax **group-interface** *ip-int-name* [**create**] [**type**]
group-interface *ip-int-name*

Context config>service>ies>subscriber-interface
config>service>vprn>subscriber-interface

Description This command creates a group interface. This interface is designed for triple-play services where multiple SAPs are part of the same subnet. A group interface may contain one or more SAPs.

The **no** form of this command removes the group interface from the subscriber interface.

Parameters *ip-int-name* — Specifies the interface name of a group interface. If the string contains special characters (#, $, spaces, and so on.), the entire string must be enclosed within double quotes.

type — Specifies the interface type.

Values bonding — Specifies to use connection bonding.
gtp — Specifies to use GTP.
lns — Specifies to use LNS.
wlangw — Specifies to use WLANGW.

bonding-parameters

Syntax **bonding-parameters**

Context config>service>vprn>sub-if>grp-if
config>service>ies>sub-if>grp-if

Description This command enables the context to configure ESM connection bonding parameters. The configuration of parameters under this context is only allowed when the group interface is created with the bonding parameter specified.

connection

Syntax **connection** *connection-id* [**create**]
no connection *connection-id*

Context config>service>vprn>sub-if>grp-if>bonding-parameters
config>service>ies>sub-if>grp-if>bonding-parameters

Description This command configures a node where per-connection parameters can be defined. The ID is used as a connection identifier for bonding whenever differentiation between connections is required.

The **no** form of this command removes the connection configuration from this bonding context, which can only be done when bonding is administratively disabled.

Parameters *connection-id* — Specifies the connection ID to be assigned to connections matching the node's parameters.

Values 1,2

service

Syntax **service** *service-id*
no service

Context config>service>vprn>sub-if>grp-if>bonding-parameters>connection
config>service>ies>sub-if>grp-if>bonding-parameters>connection

Description This command binds a specified service to this connection. ESM subscribers created under this service are eligible for bonding in this group interface and are identified by the provisioned connection ID. All connections in one bonding context must use subscriber interfaces in separate router instances.

The **no** form of this command removes the service from this bonding context, which can only be done when bonding is administratively disabled.

Parameters *service-id* — Specifies the service ID of the service containing this subscriber interface.

fpe

Syntax **fpe** *fpe-id*
no fpe

Context config>service>vprn>sub-if>grp-if>bonding-parameters
config>service>ies>sub-if>grp-if>bonding-parameters

Description This command specifies which FPE is used to provision bonding functionality. The FPE cannot be changed when there are active bonded subscribers.

The **no** form of this command disables the FPE for bonding functionality under this group interface.

multicast

Syntax **multicast**

Context config>service>vprn>sub-if>grp-if>bonding-parameters
config>service>ies>sub-if>grp-if>bonding-parameters

Description This command enables the context to configure multicast in a bonding environment.

connection

Syntax **connection** *connection-id*
connection use-incoming

Context config>service>vprn>sub-if>grp-if>bonding-parameters>mcast
config>service>ies>sub-if>grp-if>bonding-parameters>mcast

Description This command configures the connection that should be used for sending out multicast traffic in a bonding context. Traffic can either be forced to use one *connection-id* or follow the connection where the setup message (IGMP/MLD) was received first (**use-incoming**).

The **no** form of this command removes the connection configuration from this bonding context, which can only be done when bonding is administratively disabled.

Default connection use-incoming

Parameters *connection-id* — Specifies the connection ID.

Values 1, 2

shutdown

Syntax [**no**] **shutdown**

Context config>service>vprn>sub-if>grp-if>bonding-parameters
config>service>ies>sub-if>grp-if>bonding-parameters

Description The **shutdown** command administratively disables the entity. When a bonding context is shut down, all bonding subscribers are removed and no new bonding subscribers can be created in this context. The bonding configuration can be altered.

When a bonding context is placed in **no shutdown**, bonding subscribers can be created with connections in the specified subscriber interfaces. The specified FPE and connections can no longer be changed in this mode.

Default shutdown

cflowd-parameters

Syntax **cflowd-parameters**

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command enables the context to configure cflowd parameters on the interface.

sampling

Syntax **sampling** {**unicast** | **multicast**} **type** {**acl** | **interface**} [**direction** {**ingress-only** | **egress-only** | **both**}]
no sampling {**unicast** | **multicast**}

Context config>service>ies>if>cflowd-parameters
config>service>ies>sub-if>grp-if>cflowd-parameters
config>service>vprn>if>cflowd-parameters
config>service>vprn>nw-if>cflowd-parameters
config>service>vprn>sub-if>grp-if>cflowd-parameters

Description This command enables and configures the cflowd sampling behavior to collect traffic flow samples through a router for analysis.

This command can be used to configure the sampling parameters for unicast and multicast traffic separately. If sampling is not configured for either **unicast** or **multicast** traffic, then that type of traffic will not be sampled.

If cflowd is enabled without either **egress-only** or **both** specified or with the **ingress-only** keyword specified, then only ingress sampling will be enabled on the associated IP interface.

The **no** form of the command disables the associated type of traffic sampling.

Parameters **unicast | multicast** — Specifies unicast or multicast sampling.

type — Specifies the cflowd sampling type on the given virtual router interface.

Values acl — Specifies ACL cflowd analysis be applied to the given virtual router interface.

interface — Specifies interface cflowd analysis be applied to the given virtual router interface

direction — Specifies the direction of the cflowd analysis that is applied to the given virtual router interface.

Values ingress-only — Specifies an ingress only direction of the cflowd analysis be applied to the given virtual router interface.

egress-only — Specifies an egress only direction of the cflowd analysis be applied to the given virtual router interface.

both — Specifies both ingress and egress direction of the cflowd analysis be applied to the given virtual router interface.

data-trigger

Syntax **data-trigger**

Context config>service>vprn>sub-if>grp-if
config>service>ies>sub-if>grp-if

Description This command enables the context to configure data-triggered subscriber management entities.

diameter-application-policy

Syntax **diameter-application-policy** *policy-name*
no diameter-application-policy

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command configures the Diameter application policy.

The **no** form of this command reverts to the default.

Parameters *policy-name* — Specifies the Diameter application policy name.

diameter-auth-policy

Syntax **diameter-auth-policy** *name*
no diameter-auth-policy

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command is used to configure the Diameter NASREQ application policy to use for authentication.

The **no** form of this command reverts to the default.

Parameters *name* — Specifies the name of the Diameter NASREQ application policy to use for authentication.

gtp-parameters

Syntax **gtp-parameters**

Context config>service>vprn>sub-if>grp-if
config>service>ies>sub-if>grp-if

Description This command enables the context to configure GTP parameters. The configuration of parameters under this context is only allowed when the group interface is created with the GTP parameter specified.

fpe

Syntax **fpe** *fpe-id*
no fpe

Context config>service>vprn>sub-if>grp-if>gtp-parameters
config>service>ies>sub-if>grp-if>gtp-parameters

Description This command configures the FPE to be used by a group interface for extended ESM functionality such as GTP termination or bonding.

The FPE must be configured in mode sub-mgmt-extension and must be provisioned before the extended functionality becomes active.

The **no** form of this command disables the FPE for GTP functionality under this group interface.

ingress

Syntax **ingress**

Context config>service>vprn>sub-if>grp-if

Description This command enables the context to configure ingress network filter parameters for the interface.

policy-accounting

Syntax **policy-accounting** *template-name*
no policy-accounting

Context config>service>vprn>sub-if>grp-if>ingress

Description This command configures the specified policy accounting template.

The **no** form of this command disables the policy accounting template.

Parameters *template-name* — Specifies the template name up to 32 characters.

ip-mtu

Syntax **ip-mtu** *octets*
no ip-mtu

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command specifies the maximum size of IP packets on this group interface. Packets larger than this are fragmented.

The **ip-mtu** applies to all IPoE host types (dhcp, arp, static). For PPP/L2TP sessions, the **ip-mtu** is not taken into account for the MTU negotiation. The **ppp-mtu** in the **ppp-policy** should be used instead.

The **no** form of this command reverts to the default.

Parameters *octets* — Specifies the largest frame size (in octets) that this interface can handle.

Values 512 to 9000

enable-ingress-stats

Syntax [**no**] **enable-ingress-stats**

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command enables the collection of ingress interface IP stats. This command is only applicable to IP statistics, and not to uRPF statistics.

If enabled, then the following statistics are collected:

- IPv4 offered packets
- IPv4 offered octets
- IPv6 offered packets
- IPv6 offered octets

Note: Octet statistics for IPv4 and IPv6 bytes at IP interfaces include the Layer 2 frame overhead.

The **no** form of this command reverts to the default.

host-connectivity-verify

Syntax **host-connectivity-verify** [**interval** *interval*] [**action** {**remove** | **alarm**}] [**timeout** *retry-timeout*] [**retry-count** *count*] [**family** *family*]

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command enables subscriber host connectivity verification on a given SAP within a service. This tool periodically scans all known hosts (from dhcp-state) and perform UC ARP requests. The subscriber host connectivity verification maintains state (connected versus. not-connected) for all hosts.

The **no** form of this command reverts to the default.

Parameters *interval* — Specifies the interval, in minutes, which specifies the time interval which all known sources should be verified. The actual rate is then dependent on the number of known hosts and interval.

Values 1 to 6000

A zero value can be used by the SNMP agent to disable host-connectivity-verify.

action {remove | alarm} — Defines the action taken on a subscriber host connectivity verification failure for a given host. The **remove** keyword raises an alarm and removes dhcp-state and releases all allocated resources (queues, table entries and so on). DHCP-RELEASE is signaled to corresponding DHCP server. Static hosts is never removed. The **alarm** keyword raises an alarm indicating that the host is disconnected.

retry-timeout — Specifies the retry timeout.

Values 10 to 60 seconds

count — Specifies the number of retry requests.

Values 2 to 29

family — Allows the host connectivity checks to be performed for IPv4 endpoint, IPv6 endpoint or both. With family IPv6 configured, host connectivity checks is performed on the global unicast address (assigned via SLAAC or DHCPv6 IA_NA) and link-local address of a Layer 3 RG or bridged hosts. In case of SLAAC assignment, host connectivity can only be performed if the /128 is known (via downstream ND). DHCPv6 PD assigned prefixes is removed if link-local address is determined to be unreachable via host connectivity check” Reachability checks for GUA and link-local address is done simultaneously.

Values ipv4, ipv6, both

icmp

Syntax **icmp**

Context config>service>ies>if
config>service>ies>sub-if>grp-if
config>service>vprn>if

Description This command enables the context to configure Internet Control Message Protocol (ICMP) parameters on a service.

mask-reply

Syntax [**no**] **mask-reply**

Context config>service>ies>if>icmp
config>service>ies>sub-if>grp-if
config>service>vprn>if

Description This command enables responses to Internet Control Message Protocol (ICMP) mask requests on the router interface.

If a local node sends an ICMP mask request to the router interface, the **mask-reply** command configures the router interface to reply to the request.

By default, the router instance replies to mask requests.

The **no** form of this command disables replies to ICMP mask requests on the router interface.

Default mask-reply

redirects

Syntax **redirects** [*number seconds*]
no redirects

Context config>service>ies>if>icmp
config>service>ies>sub-if>grp-if>icmp
config>service>vprn>if

Description This command configures the rate for Internet Control Message Protocol (ICMP) redirect messages issued on the router interface.

When routes are not optimal on this router and another router on the same subnetwork has a better route, the router can issue an ICMP redirect to alert the sending node that a better route is available.

The **redirects** command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects is issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of redirect messages that can be issued on the interface for a given time interval.

The **no** form of this command disables the generation of icmp redirects on the router interface.

Default redirects 100 10

Parameters *number* — Specifies the maximum number of ICMP redirect messages to send. This parameter must be specified with the *second* parameter.

Values 10 to 1000

seconds — Specifies the time frame in seconds used to limit the *number* of ICMP redirect messages that can be issued.

Values 1 to 60

ttl-expired

Syntax **ttl-expired** *number seconds*
no ttl-expired

Context config>service>ies>if>icmp

Description This command configures the rate Internet Control Message Protocol. (ICMP) TTL expired messages are issued by the IP interface.
By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the limiting the rate of TTL expired messages on the router interface.

Default ttl-expired 100 10

Parameters *number* — Specifies the maximum number of ICMP TTL expired messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 to 2000

seconds — Specifies the time frame, in seconds, used to limit the number of ICMP TTL expired messages that can be issued, expressed as a decimal integer.

Values 1 to 60

unreachables

Syntax **unreachables** [*number seconds*]
no unreachables

Context config>service>ies>if>icmp

Description This command configures the rate for ICMP host and network destination unreachable messages issued on the router interface.

The **unreachables** command enables the generation of ICMP destination unreachables on the router interface. The rate at which ICMP unreachables is issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of destination unreachable messages which can be issued on the interface for a given time interval.
By default, generation of ICMP destination unreachables messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the generation of icmp destination unreachables on the router interface.

Default unreachables 100 10

Parameters *number* — Specifies the maximum number of ICMP unreachable messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 to 2000

seconds — Specifies the time frame, in seconds, used to limit the *number* of ICMP unreachable messages that can be issued.

Values 1 to 60

ipoe-linking

Syntax **ipoe-linking**

Context config>service>ies>sub-if
config>service>vprn>sub-if
config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command enables the context to configure IPoE host linking.

gratuitous-rtr-adv

Syntax [**no**] **gratuitous-rtr-adv**

Context config>service>ies>sub-if>ipoe-linking
config>service>vprn>sub-if>ipoe-linking
config>service>ies>sub-if>grp-if>ipoe-linking
config>service>vprn>sub-if>grp-if>ipoe-linking

Description This command enables the generation of unsolicited Router-advertisement on creation of v4 host.

The **no** form of this command disables **gratuitous-rtr-adv.**

Default gratuitous-rtr-adv

ipoe-session

Syntax [**no**] **ipoe-session**

Context config>service>ies>sub-if
config>service>vprn>sub-if

Description This command enables the context to configure IPoE session parameters.

shared-circuit-id

Syntax [**no**] **shared-circuit-id**

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description If configured, circuit-id in DHCPv4 Option82 is used to authenticate DHCPv6. If DHCPv6 is received before DHCPv4, it is dropped. Also, a SLAAC host is created based on DHCPv4 authentication if RADIUS returns IPv6 framed-prefix. The IPv6oE host is deleted if the linked IPv4oE host is deleted due to DHCP release or lease time-out. The linkage between IPv4 and IPv6 is based on SAP and MAC address. The sharing of circuit-id from DHCPv4 for authentication of DHCPv6 (or SLAAC) allows 7750 SR to work around lack of support for LDRA on access nodes.

The **no** form of this command reverts to the default.

qos-route-lookup

Syntax [**no**] **qos-route-lookup**

Context config>service>ies>sub-if>grp-if
config>service>ies>sub-if>grp-if>ipv6

Description This command enables QoS route lookup for the interface.

The **no** form of this command reverts to the default (strict) mode.

local-proxy-arp

Syntax [**no**] **local-proxy-arp**

Context config>service>ies>if
config>service>ies>sub-if>grp-if
config>service>vprn>if
config>service>vprn>sub-if>grp-if

Description This command enables local proxy ARP. When local proxy ARP is enabled on an IP interface, the system responds to all ARP requests for IP addresses belonging to the subnet with its own MAC address, and thus becomes the forwarding point for all traffic between hosts in that subnet. When **local-proxy-arp** is enabled, ICMP redirects on the ports associated with the service are automatically blocked.

The **no** form of this command reverts to the default.

loopback

Syntax [**no**] **loopback**

Context config>service>ies>if
config>service>vprn>if

Description This command specifies that the associated interface is a loopback interface that has no associated physical interface. As a result, the associated IES/VPRN interface cannot be bound to a SAP.

Note: Configure an IES interface as a loopback interface by issuing the **loopback** command instead of the **sap** *sap-id* command. The loopback flag cannot be set on an interface where a SAP is already defined and a SAP cannot be defined on a loopback interface.

The **no** form of this command reverts to the default.

mac

Syntax **mac** *ieee-address*
no mac

Context config>service>ies>if
config>service>ies>sub-if>grp-if

Description This command assigns a specific MAC address to a subscriber group interface.

The **no** form of this command returns the MAC address of the group interface to the default value.

Default The physical MAC address associated with the Ethernet interface that the SAP is configured on (the default MAC address assigned to the interface, assigned by the system).

Parameters *ieee-address* — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

oper-up-while-empty

Syntax [**no**] **oper-up-while-empty**

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command allows the subscriber interface to treat this group interface to be operationally enabled without any active SAPs.

This command is typically used with MSAPs where advertising the subnet prior to having a MSAP dynamically created is needed.

The **no** form of this command reverts to the default.

policy-control

Syntax **policy-control** *diameter-policy-name*
no policy-control

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command configures a policy-control policy for the interface.

The **no** form of this command reverts to the default.

Parameters *diameter-policy-name* — Specifies the name of an existing diameter policy, up to 32 characters.

proxy-arp-policy

Syntax [**no**] **proxy-arp-policy** *policy-name* [*policy-name*]

Context config>service>ies>if
config>service>ies>sub-if>grp-if
config>service>vprn>if
config>service>vprn>sub-if>grp-if

Description This command configures a proxy ARP policy for the interface.

The **no** form of this command disables the proxy ARP capability.

Parameters *policy-name* — Specifies the export route policy name. Allowed values are any string, up to 32 characters, composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

The specified name must already be defined.

sap

Syntax **sap** *sap-id* [**create**]
no sap *sap-id*

Context config>service>ies>if
config>service>ies>sub-if>grp-if
config>service>vprn>if
config>service>vprn>sub-if>grp-if

Description This command creates a Service Access Point (SAP) within a service. A SAP is a combination of port and encapsulation parameters which identifies the service access point on the interface and within the router. Each SAP must be unique.

All SAPs must be explicitly created. If no SAPs are created within a service or on an IP interface, a SAP will not exist on that object.

Enter an existing SAP without the **create** keyword to edit SAP parameters. The SAP is owned by the service in which it was created.

A SAP can only be associated with a single service. A SAP can only be defined on a port that has been configured as an access port using the **config interface** *port-type port-id* **mode access** command.

If a port is shutdown, all SAPs on that port become operationally down. When a service is shutdown, SAPs for the service are not displayed as operationally down although all traffic traversing the service is discarded. The operational state of a SAP is relative to the operational state of the port on which the SAP is defined.

Note: Configure an IES interface as a loopback interface by issuing the **loopback** command instead of the **sap** *sap-id* command. The loopback flag cannot be set on an interface where a SAP is already defined and a SAP cannot be defined on a loopback interface.

The **no** form of this command deletes the SAP with the specified port. When a SAP is deleted, all configuration parameters for the SAP are also deleted. For IES, the IP interface must be shutdown before the SAP on that interface may be removed.

Special Cases **IES —** IES SAP can be defined with Ethernet or SONET/SDH ports on the 7450 ESS. A SAP is defined within the context of an IP routed interface. Each IP interface is limited to a single SAP definition. Attempts to create a second SAP on an IP interface fails and generates an error; the original SAP is not affected.

mode

Syntax **mode** *mode*

Context config>card>mda>atm

Description This command configures the ATM MDA into a mode with the increased VC scale (16k VCs, as opposed to 8K VCs). ESM is supported only in 16K VCs mode. In 16K VCs mode, there is only one queue allocated to each VC in the ATM MDA. In 8K VCs mode, there are two queues allocated per VC.

The 16K VC mode is supported only on the 4-port oc-3/12c/STM-1/4c ATM MDA.

Changing the ATM MDA mode requires a reset of the MDA. A warning is issued asking for the confirmation before the command is executed.

Default mode max8k-vc

Parameters *mode* — Specifies the VC scale.

Values max8k-vc, max16k-vc

agg-rate

Syntax [**no**] **agg-rate**

Context config>service>ies>sub-if>grp-if>sap>egress
config>service>vprn>sub-if>grp-if>sap>egress

Description This command enables the context to configure aggregation rate parameters. This command is used to control an HQoS aggregate rate limit. It is used in conjunction with the following parameter commands: **rate**, **limit-unused-bandwidth**, **and queue-frame-based-accounting**.

When specified under a Vport, the **agg-rate**, **port-scheduler-policy** and **scheduler-policy** commands are mutually exclusive. Changing between the use of a scheduler policy and the use of an **agg-rate** or **port-scheduler-policy** involves removing the existing command and applying the new command.

The **no** form of this command disables the aggregation rate.

rate

Syntax **rate** *rate*
no rate

Context config>service>vprn>sub-if>grp-if>sap>egress>agg-rate
config>service>ies>sub-if>grp-if>sap>egress>agg-rate

Description This command defines the enforced aggregate rate for all queues associated with the agg-rate context. A rate must be specified for the **agg-rate** context to be considered to be active on the context's object (SAP, subscriber, Vport, and so on).

The **no** form of this command reverts to the default.

Parameters *rate* — Specifies the rate limit in kb/s

limit-unused-bandwidth

Syntax [**no**] **limit-unused-bandwidth**

Context config>service>ies>sub-if>grp-if>sap>egress>agg-rate
config>service>vprn>sub-if>grp-if>sap>egress>agg-rate

Description This command enables aggregate rate overrun protection.

The **no** form of this command reverts to the default.

queue-frame-based-accounting

Syntax [**no**] **queue-frame-based-accounting**

Context config>service>vprn>sub-if>grp-if>sap>egress>agg-rate
config>service>ies>sub-if>grp-if>sap>egress>agg-rate

Description This command enables frame-based accounting on all queues associated with the **agg-rate** context. Only supported on Ethernet ports. Not supported on HSMDA Ethernet ports. Packet byte offset settings are not included in the applied rate when queue frame based accounting is configured, however the offsets are applied to the statistics.

The **no** form of this command disables frame-based accounting.

vpi

Syntax **vpi** *vpi* **egress-traffic-desc** *atm-td-profile-id*
no vpi *vpi*

Context config>port>sonet-sdh>path>atm

Description This command enables the ATM VP shaper under the ATM port. The type of ATM shaper are CBR or rt/nrt-VBR as defined by the traffic descriptor. It cannot be a UBR service-type.

All VCs within the shaper degrades into a UBR type service class. For example, when a CBR type VC is associated with the shaper, it degrades into a UBR type VC. Scheduling traffic amongst VCs within the shaper is based on WRR using the weight parameter.

If the VP shaper is deleted, the VCs that were under it is restored to their original service category.

The VP shaper is statically configured and instantiated upon configuration.

A VP shaper can be seamlessly added to or removed from the active VCs in the system.

The **no** form of this command reverts to the default.

Parameters *vpi* — Specifies the VP identifier.

Values 0 to 4095

atm-td-profile-id — Specifies the ATM traffic description ID.

Values 1 to 1000

traffic-desc

Syntax **traffic-desc** *atm-td-profile-id*
no traffic-desc

Context config>service>vprn>sub-if>grp-if>sap>atm>egress
config>service>vprn>sub-if>grp-if>sap>atm>ingress
config>service>ies>sub-if>grp-if>sap>atm>egress
config>service>ies>sub-if>grp-if>sap>atm>ingress
config>subscr-mgmt>msap-policy>atm>egress
config>subscr-mgmt>msap-policy>atm>ingress

Description This command references traffic-descriptor id for VPs and VCs. The VP shaper cannot be of service-type UBR.

The **no** form of this command reverts to the default.

Default Default traffic descriptor (id=1) of UBR type.

Parameters *atm-td-profile-id* — Specifies the traffic-descriptor ID.

Values 1 to 1000

weight

Syntax **weight** *weight*
no weight

Context config>qos>atm-td-profile

Description VCs within the VP tunnel are serviced by a single scheduler assigned to each VP tunnel. VCs within the shaped VP tunnel are degraded from the originally assigned service category to a common UBR service category (default traffic descriptor). Scheduling between VCs are based on WRR with a weight parameter that can be explicitly configured in the ATM traffic descriptor profile. If weight is not specifically configured, the defaults are taken.

The explicitly configured weight parameter is honored only on the ATM MDA in the max16k-vc mode. On all other ATM capable MDAs (ASAP or ATM MDA in max8k-vc mode), the weight parameter is ignored.

The **no** form of this command reverts to the default.

Default VC degraded from CBR = weight 10

VC degraded from rt-VBR = weight 7

VC degraded from nrt-VBR = weight 5

VC degraded from UBR+ = weight 2

VC degraded from UBR = weight 1

Parameters *weight* — Specifies the weight parameter value.

Values 1 to 255

encapsulation

Syntax **encapsulation** [**aal5auto** | **aal5nlpid-ppp** | **aal5mux-ppp** | **aal5snap-bridged** | **aal5mux-bridged-eth-nofcs**]
no encapsulation

Context config>service>ies>sub-if>grp-if>sap>atm
config>service>vprn>sub-if>grp-if>sap>atm
config>service>vpls>sap>atm

Description This command is a SAP level command that either statically sets or enable dynamic detection of the encapsulation.

The **no** form of this command reverts to the default.

Parameters **aal5auto** — This option is available only in max16k-vc mode on dynamic or static SAPs. Enables automatic detection of one of the four supported encapsulation types.

aal5mux-bridged-eth-nofcs — This option already exist outside of the ESM context on regular interfaces. Within the ESM context (group-interfaces and capture SAPs), this option is available only in max16K-vc mode. The encapsulation is statically set to VC-MUX bridged Ethernet with no FCS. This is a valid encapsulation only for PPPoEoA.

aal5mux-ppp — This option is available only in max16k-vc mode on dynamic or static SAPs. The encapsulation is statically set VC-MUX PPP encapsulation. This is a valid encapsulation only for PPPoA.

aal5nlpid-ppp — dynamic or static SAPs. The encapsulation is statically set to NLPID (LLC) PPP encapsulation. This is a valid encapsulation only for PPPoA.

aal5snap-bridged — This option already exist outside of the ESM context on regular interfaces. Within the ESM context (group-interfaces and capture SAPs), this option is available only in max16k-vc mode. The encapsulation is statically set to bridged Ethernet with or without FCS. Both PIDs (0x 00-01 and 0x 00-07) are accepted on ingress and use this to determine whether to strip four bytes from the end of the encapsulated Ethernet frame. The inner FCS is not checked. This is a valid encapsulation only for PPPoEoA.

Note: On ATM, frames with Ethernet FCS or without FCS, are accepted but only frames with no Ethernet FCS are sent.

def-inter-dest-id

Syntax **def-inter-dest-id string** *interest-string*
def-inter-dest-id {**use-top-q** | **use-vpi**}
no def-inter-dest-id

Context config>service>ies>sub-if>grp-if>sap>sub-sla-mgmt
config>service>vprn>sub-if>grp-if>sap>sub-sla-mgmt
config>subscr-mgmt>msap-policy>sub-sla-mgmt

Description This command associates the Vport with the subscriber. The association method depends on the configured option.

The **no** form of this command reverts to the default.

Parameters *string* — A RADIUS VSA (Alc-Int-Dest-Id-Str, type 28) obtained during the subscriber authentication phase contains the destination string name that is matched against the string defined under the Vport. In this fashion the subscriber host is associated with the corresponding Vport.

Alternatively, the destination string can be defined in LUDB.

use-top-q — This is applicable only to Ethernet ports.

use-vpi — VP Identifier (VPI) is used to make the association between the subscriber and the Vport automatically.

The control plane is aware of the VPI during the session initiation phase. This VPI is used to make the association between the host and the Vport with the same name (VPI number).

Note: In this case the Vport name under the **configure>port>sonet-sdh>path>access>egress** context must be the VPI number.

pppoe-user-db

Syntax **pppoe-user-db** *ludb-name*
no pppoe-user-db

Context config>services>vpls>sap

Description This command enables LUDB authentication on capture SAPs for PPPoE(oA) clients. If this command is configured along with the **authentication-policy** command (RADIUS authentication), then the authentication-policy command takes precedence.

Optionally, with a separate command (ppp-user-db) PPPoA clients can be authenticated under a separate LUDB.

The **no** form of this command reverts to the default.

Parameters *ludb-name* — Specifies the name of the local user database up to 32 characters.

ppp-user-db

Syntax **ppp-user-db** *ludb-name*
no ppp-user-db

Context config>services>vpls>sap

Description This command enables LUDB authentication on capture SAPs for PPPoA clients. If this command is configured along with the authentication-policy command (RADIUS authentication), then the authentication-policy command takes precedence.

Optionally, with a separate command (**pppoe-user-db**) PPPoE(oA) clients can be authenticated under a separate LUDB.

The **no** form of this command reverts to the default.

Parameters *ludb-name* — Specifies the name of local user database up to 256 characters.

ppp-policy

Syntax **ppp-policy** *ppp-pol-name*
no ppp-policy

Context config>services>vpls>sap

Description This command references a ppp-policy that defines session parameters (ppp-mtu, authentication options, and so on) during the session initiation phase. Normally, the PPP policy is referenced under the group-interface hierarchy. But with capture SAP is it not known at the session initiation phase to which group-interface the session belongs. This is why, with the capture SAP, the ppp-policy must be referenced directly under the capture SAP. The ppp-policy referenced under the **group-interface** must be the same as the ppp-policy referenced under the capture SAP. Otherwise the session will not come up.

The **no** form of this command reverts to the default.

Parameters *ppp-pol-name* — Specifies the PPP policy name up to 32 characters.

pppoe-policy

Syntax **pppoe-policy** *pppoe-policy-name*
no pppoe-policy

Context config>services>vpls>sap

Description This command references a pppoe-policy that defines session parameters (ppp-mtu, authentication options, and so on) during the session initiation phase. Normally, the PPPoE policy is referenced under the **group-interface** hierarchy. But with capture SAP is it not known at the session initiation phase to which group-interface the session belongs. This is why, with the capture SAP, the ppp-policy must be referenced directly under the capture SAP. The pppoe-policy referenced under the group-interface must be the same as the pppoe-policy referenced under the capture SAP. Otherwise the session will not come up.

The **no** form of this command reverts to the default.

Parameters *pppoe-policy-name* — Specifies the pppoe-policy name up to 32 characters.

pppoe-python-policy

Syntax **pppoe-python-policy** *policy-name*
no pppoe-python-policy

Context config>service>vpls>sap

Description This command specified the Python policy for PPPoE packets sent/received on the capture SAP.

The **no** form of this command removes the policy name from the configuration.

Parameters *policy-name* — Specifies an existing Python policy name

rtr-solicit-user-db

Syntax **rtr-solicit-user-db** *local-user-db*
no rtr-solicit-user-db

Context config>service>vpls>sap

Description This command enabled access to LUDB for SLAAC hosts under the capture SAP. The name of this ludb must match the name of ludb configured under the **configure>service>vprn/ies>sub-if>group-if>ipv6>router-solicit** hierarchy.

The **no** form of this command reverts to the default.

Parameters *local-user-db* — Specifies the name of the local-user-database up to 32 characters.

track-srrp

Syntax **track-srrp** *srrp-id*
no track-srrp

Context config>service>vpls>sap

Description This is a capture SAP level command. This command is important in PPPoE deployments with MSAPs. PPPoE operation requires that the MAC address learned by the client at the very beginning of the session negotiation phase remains unchanged for the lifetime of the session (RFC 2516). This command ensures that the virtual MAC address used during the PPPoE session negotiation phase on the capture SAP is the same virtual MAC address that is used by the SRRP on the group interface on which the session is established. Therefore, it is mandated that the SRRP instance (and implicitly the group-interface) where the session belongs to is known in advance. If the group interface name for the session is returned by the RADIUS, it must be ensured that this group interface is the one on which the tracked SRRP instance is configured. PPPoE sessions on the same capture SAP cannot be shared across multiple group interfaces, but instead they all must belong to a single group interface that is known in advance.

The same restrictions apply to IPoE clients in MC Redundancy scenario if they are to be supported concurrently on the same capture SAP as PPPoE.

The supported capture SAP syntax is this:

sap <port-id>:X.* capture-sap

The capture SAP syntax that is not supported is this:

sap <port-id>:*.* capture-sap

The **no** form of this command removes the SRRP ID from this configuration.

Parameters *srrp-id* — Specifies the SRRP instance number.

Values 1 to 4294967295

vc-range

Syntax **vc-range** *num* **vpi-range** *vpi-range* **vci-range** *vci-range*
no vc-range num

Context config>services>vpls>sap>atm

Description This command configures ATM Virtual Circuit ranges and is supported only in max16k-vc ATM MDA mode. An ATM MDA supports a number of passive (or listening) VCs, of which a subset can be simultaneously active.

The **no** form of this command reverts to the default.

Parameters *num* — Specifies the VC range.

Values 1 to 5 (Five ranges are supported to accommodate non-contiguous ranges of VPI/VCI pairs.

vci-range *vci-range* — Specifies the VCI range.

Values 1, 2, 5 to 65535 (Contiguous VCI ranges in the form of x-y.)

vpi-range vpi-range. — Specifies the VPI range.

Values 0 to 255 for UNI
0 to 4095 for NNI
(Contiguous VPI range in the form of x-y)

local-address-assignment

Syntax **local-address-assignment**

Context config>service>ies>sub-if
config>service>vprn>sub-if
config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command enables the context to configure local address assignment parameters.

ipv6

Syntax **ipv6**

Context config>service>ies>sub-if
config>service>vprn>sub-if

config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command enables the context to configure IPv6 parameters.

client-application

Syntax **client-application** [**ppp-v4**] [**ipoe-v4**]
no client-application

Context config>service>ies>sub-if>lcl-addr-assign
config>service>vprn>sub-if>lcl-addr-assign
config>service>ies>sub-if>grp-if>lcl-addr-assign
config>service>vprn>sub-if>grp-if>lcl-addr-assign

Description This command enables Local Address Assignment selectively for PPP or IPoE clients. IPoE is only supported for certain scenarios where no DHCP is used.

When enabled, a pool can be shared between non-DHCP clients and regular DHCP clients in the same DHCP server and pool.

The **no** form of this command disables Local Address Assignment for any protocol.

client-application

Syntax **client-application** [**ppp-slaac**] [**ipoe-wan**] [**ipoe-slaac**]
no client-application

Context config>service>vprn>sub-if>grp-if>lcl-addr-assign>ipv6
config>service>vprn>sub-if>grp-if>lcl-addr-assign>ipv6

Description This command defines the client application that uses the local address server to perform address assignment. This feature is relies on RADIUS or local-user-database to return a pool name. The pool name is matched again the pools defined in the local-dhcp6-server configuration. The name of the local-dhcp6-server must also be provisioned.

The **no** form of this command reverts to the default.

Parameters **ppp-slaac** — Indicates using the local DHCPv6 prefix pool to assign SLAAC prefixes for hosts. The pool name where the prefixes are used for SLAAC prefix assignment are obtained from RADIUS or local-user-database during the authentication process. The RADIUS attribute Alc-slaac-ipv6-pool is used to indicate the SLAAC pool name for PPPoE hosts.

ipoe-wan — Indicates using the local DHCPv6 pool for IA_NA address assignment and a static pre-defined prefixes for IA_PD. Both the IA_NA pool name and the IA_PD static framed-prefix are either obtained from RADIUS or LUDB during authentication. With RADIUS, it must return both IA_NA Framed-IPv6-Pool and IA_PD Delegated-IPv6-Prefix after a successful authentication. With LUDB, it must have ipv6-wan-address-pool and ipv6-delegated-prefix populated. This feature is specific to this use case and is not required for other combinations of DHCPv6 assignments such as IA_NA and IA_PD address assignment through RADIUS or LUDB.

ipoe-slaac — Indicates using the local DHCPv6 prefix pool to assign SLAAC prefixes for hosts. The pool name where the prefixes are used for SLAAC prefix assignment are obtained from RADIUS or local-user-database during the authentication process. The RADIUS attribute Alc-slaac-ipv6-pool is used to indicate the SLAAC pool name for PPPoE hosts.

unnumbered

Syntax **unnumbered** [*ip-int-name* | *ip-address*]
no unnumbered

Context config>service>vprn>sub-if
config>service>ies>sub-if

Description This command can be configured only for subscriber interfaces that do not have an IPv4 address explicitly configured and is therefore operationally in a DOWN state. By configuring this command, the subscriber interface borrows the IPv4 address from the referenced interface (directly or indirectly via IP address) that must be operationally UP and located in the same routing instance as the subscriber interface. This allows the subscriber interface to becomes operationally UP and consequently allow forwarding of the subscriber traffic.

Such interface is referred as unnumbered interface, since it does not have explicitly configured a unique IP address. Subscriber hosts under the unnumbered subscriber interface are installed in the fib as /32 hosts.

Without this command the subscriber interface is operationally DOWN and subscriber-host instantiation is not possible.

This command is mutually exclusive with the allow-unmatched-subnets command under the same CLI hierarchy.

The operation of IPv6 host is not affected by this command.

The **no** form of this command reverts to the default.

Parameters *ip-int-name* — Specifies the interface name from which an IPv4 address is borrowed.

ip-address — Specifies the IP address from an optionally up interface that is used for subscriber interface.

default-pool

Syntax **default-pool** *pool-name* [**secondary** *pool-name*]
no default-pool

Context config>service>ies>sub-if>lcl-addr-assign
config>service>vprn>sub-if>lcl-addr-assign
config>service>ies>sub-if>grp-if>lcl-addr-assign
config>service>vprn>sub-if>grp-if>lcl-addr-assign

Description This command references a default DHCP address pool for local PPPoX pool management in case that the pool-name is not returned via RADIUS or LUDB.

The **no** form of this command reverts to the default.

Parameters *pool-name* — Specifies the name of the local router DHCP server pool.

server

Syntax **server** *server-name*
no server

Context config>service>ies>sub-if>lcl-addr-assign
config>service>vprn>sub-if>lcl-addr-assign
config>service>ies>sub-if>grp-if>lcl-addr-assign
config>service>vprn>sub-if>grp-if>lcl-addr-assign

Description This command designates a local router DHCPv4 server for local pools management where IPv4 addresses for PPPoXv4 clients is allocated without the need for the internal router DHCP relay-agent. Those addresses are tied to PPPoX sessions and they are de-allocated when the PPPoX session is terminated.

Parameters *server-name* — The name of the local 7450 ESS or 7750 SR DHCP server.

server

Syntax **server** *server-name*
no server

Context config>service>ies>sub-if>grp-if>lcl-addr-assign>ipv6
config>service>vprn>sub-if>grp-if>lcl-addr-assign>ipv6

Description This command designates a local router DHCPv6 server for local pools management where IPv6 prefixes or address for PPPoXv6 clients or IPoEv6 clients are allocated without the need for the internal router DHCP relay-agent. Those addresses are tied to PPPoX or IPoE sessions and they are de-allocated when the PPPoX or IPoE session is terminated.

Parameters *server-name* — The name of the local router DHCPv6 server.

private-retail-subnets

Syntax [**no**] **private-retail-subnets**

Context config>service>vprn>sub-if

Description This command controls the export of retail subnets and prefixes to the wholesale forwarding service. When this attribute is configured, subnets and prefixes configured on the retail subscriber interface are no longer exported to the associated wholesale VPRN and remain private to the retail VPRN. This is useful in a IPoE or PPPoE business service context, as it allows retail services to use overlapping IP address spaces even if those services are associated with the same wholesale service. IPoE and PPPoE sessions are actually terminated in the retail service although their traffic transits on a SAP belonging to the wholesale service.

Configuring private retail subnets is not supported for IPv4 static hosts and ARP hosts. If PPPoE sessions need to coexist with IPv4 host types, then this attribute should not be configured on the retail subscriber interface.

This command fails if the subscriber interface is not associated with a wholesale service.

If the retail VPRN is of the type **hub**, this attribute is mandatory. In this case, private retail subnets are enabled by default and cannot be unconfigured.

The **no** form of this command disables overlapping IP addresses between different retailers referring to this interface.

qos-route-lookup

Syntax [**no**] **qos-route-lookup**

Context config>service>ies>sub-if>grp-if
config>service>ies>sub-if>grp-if>ipv6

Description This command enables or disables QoS route lookup for the interface.

The **no** form of this command reverts to the default.

nd

Syntax **nd**

Context config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>grp-if>ipv6

Description This command enables the context to configure neighbor discovery (ND) parameters.

dad-snooping

Syntax [**no**] **dad-snooping**

Context config>service>ies>sub-if>grp-if>ipv6>nd
config>service>vprn>sub-if>grp-if>ipv6>nd

Description This command allows the router to populate the neighbor discovery table through snooping subscribers' duplicate address detection messages.

The **no** form of this command reverts to the default.

neighbor-limit

Syntax **neighbor-limit** [*value*]
no neighbor-limit

Context config>service>ies>sub-if>grp-if>ipv6>nd
config>service>vprn>sub-if>grp-if>ipv6>nd

Description This command configures the maximum number of neighbors learned for a single host by doing neighbor discovery.

The **no** form of this command reverts to the default.

Default neighbor-limit 1

Parameters *value* — Specifies the maximum number of neighbors learned.

Values 1 to 8

dns-options

Syntax [**no**] **dns-options**

Context config>router>router-advert
config>router>router-advert>if
config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6>rtr-adv
config>service>vprn>router-advert
config>service>vprn>router-advert>if
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv

Description This command enables the context to configure IPv6 DNS options for SLAAC hosts.

include-dns

Syntax [**no**] **include-dns**

Context config>router>rtr-adv>if>dns-opt
config>service>ies>sub-if>ipv6>rtr-adv>dns-opt
config>service>ies>sub-if>grp-if>ipv6>rtr-adv>dns-opt
config>service>vprn>rtr-adv>if>dns-opt
config>service>vprn>sub-if>ipv6>rtr-adv>dns-opt
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv>dns-opt

Description This command specifies to include the Recursive DNS Server (RDNSS) Option as defined in RFC 6106 in IPv6 Router Advertisements for DNS name resolution of IPv6 SLAAC hosts

The **no** form of this command reverts to the default.

rdnss-lifetime

Syntax **rdnss-lifetime** *seconds*
rdnss-lifetime infinite
no rdnss-lifetime

Context config>router>rtr-adv>dns-opt
config>router>rtr-adv>if>dns-opt
config>service>ies>sub-if>ipv6>rtr-adv>dns-opt
config>service>ies>sub-if>grp-if>ipv6>rtr-adv>dns-opt
config>service>vprn>rtr-adv>dns-opt
config>service>vprn>rtr-adv>if>dns-opt
config>service>vprn>sub-if>ipv6>rtr-adv>dns-opt
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv>dns-opt

Description This command configures the maximum time that the RDNSS address may be used for name resolution.

The **no** form of this command reverts to the default.

Default rdnss-lifetime 3600

Parameters *seconds* — Specifies the time, in seconds, that the RDNSS address is valid for this route

Values 900 to 3600

infinite — The RDNSS address can be used permanently.

force-mcast

Syntax **force-mcast** [**ip**] [**mac**]
no force-mcast

Context config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6>rtr-adv
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv

Description This command configures the multicast router advertisements on this interface, either IP or MAC.

The **no** form of this command reverts to the default.

managed-configuration

Syntax [**no**] **managed-configuration**

Context config>router>router-advert>if
config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6>rtr-adv
config>service>vprn>router-advert>if
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv

Description This command sets or resets managed address configuration flag for this group-interface.

The **no** form of this command reverts to the default.

max-advertisement

Syntax **max-advertisement** *seconds*
no max-advertisement

Context config>router>router-advert>if
config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6>rtr-adv
config>service>vprn>router-advert>if
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv

Description This command specifies the maximum time allowed between sending unsolicited router advertisements from this interface.

The **no** form of this command reverts to the default.

Default max-advertisement 1800

Parameters *seconds* — Specifies the maximum advertisement interval, in seconds, for this group-interface.

Values 900 to 11800

min-advertisement

Syntax **min-advertisement** *seconds*
no min-advertisement

Context config>router>router-advert>if
config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6>rtr-adv
config>service>vprn>router-advert>if
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv

Description This command specifies the minimum time allowed between sending unsolicited router advertisements from this interface.

The **no** form of this command reverts to the default.

Default min-advertisement 900

Parameters *seconds* — Specifies the minimum advertisement interval, in seconds, for this group-interface.

Values 900 to 1350

mtu

Syntax **mtu** *bytes*
no mtu

Context config>router>router-advert>if
config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6>rtr-adv
config>service>vprn>router-advert>if
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv

Description This command specifies the value to be placed in link MTU options sent by the router on this interface.

The **no** form of this command reverts to the default.

Parameters *bytes* — Sets the advertised MTU value in bytes for this group-interface.

Values 1280 to 9212

other-stateful-configuration

Syntax [**no**] **other-stateful-configuration**

Context config>router>router-advert>if
config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6>rtr-adv
config>service>vprn>router-advert>if
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv

Description This command sets and resets the other-stateful-configuration flag for this group-interface.

The **no** form of this command reverts to the default.

prefix-options

Syntax [**no**] **prefix-options**

Context config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6>rtr-adv
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv

Description This command enables the context to configure prefix options for this group-interface.

autonomous

Syntax [**no**] **autonomous**

Context config>service>ies>sub-if>ipv6>rtr-adv>pfx-opt
config>service>ies>sub-if>grp-if>ipv6>rtr-adv>pfx-opt
config>service>vprn>sub-if>ipv6>rtr-adv>pfx-op
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv>pfx-opt

Description This command enables the option that determines whether or not the prefix can be used for stateless address autoconfiguration.

The **no** form of this command disables the option.

on-link

Syntax [**no**] **on-link**

Context config>router>router-advert>if
config>service>ies>sub-if>ipv6>rtr-adv>pfx-opt
config>service>ies>sub-if>grp-if>ipv6>rtr-adv>pfx-opt
config>service>vprn>router-advert>if
config>service>vprn>sub-if>ipv6>rtr-adv>pfx-opt
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv>pfx-opt

Description This command specifies whether the prefix is assigned to an interface on the specified link.

The **no** form of this command reverts to the default.

Default on-link

preferred-lifetime

Syntax **preferred-lifetime** *seconds*
preferred-lifetime infinite
no preferred-lifetime

Context config>service>ies>sub-if>ipv6>rtr-adv>pfx-opt
config>service>ies>sub-if>grp-if>ipv6>rtr-adv>pfx-opt
config>service>vprn>sub-if>ipv6>rtr-adv>pfx-opt
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv>pfx-opt

Description This command specifies the remaining time for this prefix to be preferred, thus time until deprecation.

The **no** form of this command reverts to the default.

Default preferred-lifetime 3600

Parameters *seconds* — Specifies the time, in seconds, for the prefix to remain preferred on this interface in seconds.

Values 0 to 4294967295

infinite — Specifies that the remaining time will never expire.

reachable-time

Syntax **reachable-time** *milli-seconds*
no reachable-time

Context config>router>router-advert>if
config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6>rtr-adv
config>service>vprn>router-advert>if
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv

Description This command configures the value to be placed in the reachable time field in router advertisement messages sent from this interface.

The **no** form of this command reverts to the default.

Default reachable-time 0

Parameters *milli-seconds* — Specifies the reachable time, in seconds, for advertisements from this interface.

Values 0 to 3600000

valid-lifetime

Syntax **valid-lifetime** *seconds*
valid-lifetime infinite
no valid-lifetime

Context config>service>ies>if>ipv6>dhcp6>pfx-delegate>prefix

Description This command configures the time, in seconds, that the prefix is valid.

The **no** form of this command reverts to the default value.

Default valid-lifetime 2592000 (30 days)

Parameters *seconds* — Specifies the time, in seconds, that this prefix remains valid.

Values 1 to 4294967294

infinite — Specifies that this prefix remains valid infinitely. The value 4294967295 is interpreted as infinite.

valid-lifetime

Syntax **valid-lifetime** *seconds*
valid-lifetime infinite
no valid-lifetime

Context config>service>ies>sub-if>ipv6>rtr-adv>pfx-opt
config>service>ies>sub-if>grp-if>ipv6>rtr-adv>pfx-opt
config>service>vprn>sub-if>ipv6>rtr-adv>pfx-opt
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv>pfx-opt

Description This command specifies the remaining time for this prefix to be valid for the purpose of on-link determination.

The **no** form of this command reverts to the default.

Default valid-lifetime 86400

Parameters *seconds* — Specifies the time for the prefix to remain valid on this interface in seconds.

Values 0 to 4294967295

infinite — Specifies that the remaining time will never expire. The value 4294967295 is interpreted as infinite.

retransmit-time

Syntax **retransmit-time** *milli-seconds*
no retransmit-time

Context config>router>router-advert>if
config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6>rtr-adv
config>service>vprn>router-advert>if
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv

Description This command configures the value to be placed in the retransmit timer field in router advertisements sent from this interface.

The **no** form of this command reverts to the default.

Default retransmit-time 0

Parameters *milli-seconds* — Specifies the retransmit time, in milli-seconds, for advertisement from this group-interface.

Values 0 to 1800000

router-lifetime

Syntax **router-lifetime** *seconds*
router-lifetime no-default-router
no router-lifetime

Context config>router>router-advert>if
config>service>ies>sub-if>ipv6>rtr-adv
config>service>ies>sub-if>grp-if>ipv6>rtr-adv
config>service>vprn>router-advert>if
config>service>vprn>sub-if>ipv6>rtr-adv
config>service>vprn>sub-if>grp-if>ipv6>rtr-adv

Description This command configures the value to be placed in the router lifetime field of router advertisements sent from this interface.

The **no** form of this command reverts to the default.

Default router-lifetime 4500

Parameters *seconds* — Specifies the router lifetime in seconds for this group-interface.

Values 2700 to 9000

no-default-router — Specifies that the router is not to be used as a default router.

inactivity-timer

Syntax **inactivity-timer** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
inactivity-timer infinite
no inactivity-timer

Context config>service>ies>sub-if>ipv6>rtr-sol
config>service>ies>sub-if>grp-if>ipv6>rtr-sol
config>service>vprn>sub-if>ipv6>rtr-sol
config>service>ies>sub-if>grp-if>ipv6>rtr-sol

Description This command specifies the time before an inactive host is removed.

The **no** form of this command reverts to the default.

Parameters **infinite** — Specifies that the idle host is never removed.

days — Specifies that the idle host is removed if idle within the number of specified days.

hours — Specifies that the idle host is removed if idle within the number of specified hours.

minutes — Specifies that the idle host is removed if idle within the number of specified minutes.

seconds — Specifies that the idle host is removed if idle within the number of specified seconds.

min-auth-interval

Syntax **min-auth-interval** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no min-auth-interval

Context config>service>ies>sub-if>grp-if>ipv6>rtr-sol
config>service>ies>sub-if>grp-if>ipv6>rtr-sol

Description This command specify the minimum interval between two consecutive authentication attempts from the same host.

The **no** form of this command reverts to the default.

Parameters *days* — Specifies the number of days that a user must wait for the next authentication attempt.

hours — Specifies the number of hours that a user must wait for the next authentication attempt.

minutes — Specifies the number of minutes that a user must wait for the next authentication attempt.

seconds — Specifies the number of seconds that a user must wait for the next authentication attempt.

user-db

Syntax [**no**] **user-db**

Context config>service>ies>sub-if>grp-if>ipv6>rtr-sol
config>service>ies>sub-if>grp-if>ipv6>rtr-sol

Description This command enables the use of the local-user-database for authentication.

The **no** form of this command reverts to the default.

Parameters *local-user-db-name* — Specifies the name of the local-user-database to authenticate the router-solicit. The local-user-database can also return a static prefix or a pool name for address assignment.

shutdown

Syntax [**no**] **shutdown**

Context config>service>ies>sub-if>grp-if>ipv6>rtr-sol
config>service>ies>sub-if>grp-if>ipv6>rtr-sol

Description This command enables SLAAC triggered host creation.

The **no** form of this command disables SLAAC triggered host creation.

Default no shutdown

icmp

Syntax **icmp**

Context config>service>ies>if
config>service>ies>sub-if>grp-if
config>service>vprn>if

Description This command enables the context to configure Internet Control Message Protocol (ICMP) parameters on a service.

mask-reply

Syntax [**no**] **mask-reply**

Context config>service>ies>if>icmp
config>service>ies>sub-if>grp-if
config>service>vprn>if

Description This command enables responses to Internet Control Message Protocol (ICMP) mask requests on the router interface.

If a local node sends an ICMP mask request to the router interface, the **mask-reply** command configures the router interface to reply to the request.

By default, the router instance replies to mask requests.

The **no** form of this command disables replies to ICMP mask requests on the router interface.

redirects

Syntax **redirects** [*number seconds*]
no redirects

Context config>service>ies>if>icmp
config>service>ies>sub-if>grp-if>icmp
config>service>vprn>if

Description This command configures the rate for Internet Control Message Protocol (ICMP) redirect messages issued on the router interface.

When routes are not optimal on this router and another router on the same subnetwork has a better route, the router can issue an ICMP redirect to alert the sending node that a better route is available.

The **redirects** command enables the generation of ICMP redirects on the router interface. The rate at which ICMP redirects is issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of redirect messages that can be issued on the interface for a given time interval.

The **no** form of this command disables the generation of icmp redirects on the router interface.

Default redirects 100 10

Parameters *number* — Specifies the maximum number of ICMP redirect messages to send. This parameter must be specified with the *second* parameter.

Values 10 to 1000

seconds — Specifies the time frame in seconds used to limit the *number* of ICMP redirect messages that can be issued.

Values 1 to 60

ttl-expired

Syntax **ttl-expired** *number seconds*

no ttl-expired

Context config>service>ies>if>icmp

Description This command configures the rate Internet Control Message Protocol. (ICMP) TTL expired messages are issued by the IP interface.
By default, generation of ICMP TTL expired messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the limiting the rate of TTL expired messages on the router interface.

Default ttl-expired 100 10

Parameters *number* — Specifies the maximum number of ICMP TTL expired messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 to 2000

seconds — Specifies the time frame, in seconds, used to limit the number of ICMP TTL expired messages that can be issued, expressed as a decimal integer.

Values 1 to 60

unreachables

Syntax **unreachables** [*number seconds*]
no unreachables

Context config>service>ies>if>icmp

Description This command configures the rate for ICMP host and network destination unreachable messages issued on the router interface.

The **unreachables** command enables the generation of ICMP destination unreachables on the router interface. The rate at which ICMP unreachables is issued can be controlled with the optional *number* and *seconds* parameters by indicating the maximum number of destination unreachable messages which can be issued on the interface for a given time interval.
By default, generation of ICMP destination unreachables messages is enabled at a maximum rate of 100 per 10 second time interval.

The **no** form of this command disables the generation of icmp destination unreachables on the router interface.

Default unreachables 100 10

Parameters *number* — Specifies the maximum number of ICMP unreachable messages to send. This parameter must be specified with the *seconds* parameter.

Values 10 to 2000

seconds — Specifies the time frame, in seconds, used to limit the *number* of ICMP unreachable messages that can be issued.

Values 1 to 60

Layer 3 Subscriber Interfaces SAP Commands

sap

Syntax **sap** *sap-id* [**create**]
no sap *sap-id*

Context config>service>vprn>sub-if>grp-if

Description This command creates a SAP for the interface.

The **no** form of this command removes the SAP.

Parameters *sap-id* — Specifies the SAP ID.

accounting-policy

Syntax **accounting-policy** *acct-policy-id*
no accounting-policy

Context config>service>ies>sub-if>grp-if>sap
config>service>vprn>if>sap
config>service>vprn>if>spoke-sdp
config>service>vprn>sub-if>grp-if>sap

Description This command specifies the policy to use to collect accounting statistics on a subscriber profile.

A maximum of one accounting policy can be associated with a profile at one time.

The **no** form of this command removes the accounting policy association.

Parameters *acct-policy-id* — Specifies an existing accounting policy ID as configured in the **config>log>accounting-policy** context.

Values 1 to 99

anti-spoof

Syntax **anti-spoof** {**ip** | **ip-mac** | **nh-mac**}

no anti-spoof

Context config>service>ies>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>sap
config>subscr-mgmt>msap-policy

Description This command configures the anti-spoof type of the MSAP.

The type of anti-spoof filtering defines what information in the incoming packet is used to generate the criteria to lookup an entry in the anti-spoof filter table. The type parameter (**ip**, **ip-mac**) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled.

The **no** form of this command reverts to the default.

Note: For IES and VPRN subscriber group interfaces, setting no anti-spoof sets the default anti-spoofing type which is **ip-mac**.

Parameters **ip** — Configures SAP anti-spoof filtering to use only the source IP address in its lookup. If a static host exists on the SAP without an IP address specified, the anti-spoof type **ip** command fails.

Note: This parameter is not applicable in the **config>subscr-mgmt>msap-policy** context.

ip-mac — Configures SAP anti-spoof filtering to use both the source IP address and the source MAC address in its lookup. The anti-spoof type **ip-mac** command fails if the default anti-spoof filter type of the SAP is **ip-mac** and the default is not overridden, or if the SAP does not support Ethernet encapsulation.

nh-mac — Indicates that the ingress anti-spoof is based on the source MAC and egress anti-spoof is based on the nh-ip-address .

app-profile

Syntax **app-profile** *app-profile-name*
no app-profile

Context config>service>vprn>if>sap
config>service>vprn>sub-if>grp-if>sap

Description This command configures the application profile name.

The **no** form of this command reverts to the default.

Parameters *app-profile-name* — Specifies an existing application profile name configured in the **config>app-assure>group>policy** context.

collect-stats

Syntax [**no**] **collect-stats**

Context config>service>ies>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>sap

Description When enabled, the agent collects non-RADIUS accounting statistics on a subscriber profile.

When the **no collect-stats** command is issued the statistics are still accumulated by the IOM cards. However, the CPU does not obtain the results and write them to the billing file. If a subsequent **col-lect-stats** command is issued then the counters written to the billing file include all the traffic while the **no collect-stats** command was in effect.

The **no** form of this command removes the value from the configuration.

Default collect-stats

default-host

Syntax **default-host** *ipv4-prefix/mask* | *ipv6-prefix/prefix-length* **next-hop** *ipv4-address* | *ipv6-address*
no default-host *ipv4-prefix/mask* | *ipv6-prefix/prefix-length*

Context config>service>ies>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>sap

Description This command configures the default-host. More than one default host can be configured per SAP.

The **no** form of this command removes the values from the configuration.

Parameters *ipv64prefix/prefix-length* — Specifies an IPv4 prefix and prefix length.

Values

ipv4-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D
prefix-length	- [0 to 128]

ipv6-prefix/prefix-length — Specifies an IPv6 prefix and prefix length.

Values

ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D
prefix-length	- [0 to 128]

next-hop — Assigns the next hop IP address.

cpu-protection

Syntax **cpu-protection** *policy-id* [**mac-monitoring**] | [**eth-cfm-monitoring** [**aggregate**] [**car**]]
no cpu-protection

Context config>service>ies>sub-if>grp-if>sap

Description This command assigns an existing CPU protection policy to the associated group interface. The CPU protection policies are configured in the **config>sys>security>cpu-protection>policy** *cpu-protection-policy-id* context.

If no CPU-Protection policy is assigned to a group interface SAP, then the default policy is used to limit the overall-rate. The default policy is policy number 254 for access interfaces and 255 for network interfaces.

The **no** form of this command removes the association of the CPU protection policy from the associated interface and reverts to the default policy values.

Default cpu-protection 254 (for access interfaces)

cpu-protection 255 (for network interfaces)

The configuration of no cpu-protection returns the interface/SAP to the default policies as shown above.

Parameters *policy-id* — Specifies an existing CPU protection policy.

Values 1 to 255

mac-monitoring — Specifies to enable MAC monitoring.

eth-cfm-monitoring — Specifies to enable Ethernet Connectivity Fault Management monitoring.

aggregate — Applies the rate limit to the sum of the per peer packet rates.

car — Applies the Committed Access Rate (CAR). This keyword causes Eth-CFM packets to be ignored when enforcing the overall-rate.

egress

Syntax **egress**

Context config>service>ies>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>sap

Description This command enables the context to configure egress SAP Quality of Service (QoS) policies and filter policies.

If no sap-egress QoS policy is defined, the system default sap-egress QoS policy is used for egress processing. If no egress filter is defined, no filtering is performed.

agg-rate

Syntax [**no**] **agg-rate**

Context config>service>vprn>sub-if>grp-if>sap>egress

Description This command configures the aggregate rate limit.

The **no** form of this command removes the rate

filter

Syntax **filter ip** *ip-filter-id*
filter
no filter [**ip** *ip-filter-id*]
no filter

Context config>service>ies>sub-if>grp-if>sap>egress
config>service>ies>sub-if>grp-if>sap>ingress
config>service>vprn>sub-if>grp-if>sap>egress
config>service>vprn>sub-if>grp-if>sap>ingress

Description This command associates an IP filter policy with an ingress or egress Service Access Point (SAP). Filter policies control the forwarding and dropping of packets based on the matching criteria.

MAC filters are only allowed on Epipe and Virtual Private LAN Service (VPLS) SAPs.

The **filter** command is used to associate a filter policy with a specified *ip-filter-id* with an ingress or egress SAP. The filter policy must already be defined before the **filter** command is executed. If the filter policy does not exist, the operation fails and an error message returned.

In general, filters applied to SAPs (ingress or egress) apply to all packets on the SAP. One exception is non-IP packets are not applied to the match criteria, so the default action in the filter policy applies to these packets.

The **no** form of this command removes any configured filter ID association with the SAP. The filter ID itself is not removed from the system unless the scope of the created filter is set to local. To avoid deletion of the filter ID and only break the association with the service object, use the **scope** command within the filter definition to change the scope to **local** or **global**. The default scope of a filter is **local**.

Special Cases **IES —** Only IP filters are supported on an IES IP interface, and the filters only apply to routed traffic.

Parameters **ip** — Keyword indicating the filter policy is an IP filter.

ip-filter-id — Specifies the ID for the IP filter policy. Allowed values are an integer in the range of 1 and 65535 that corresponds to a previously created IP filter policy in the **configure>filter** context.

match-qinq-dot1p

Syntax **match-qinq-dot1p** {**top** | **bottom**}
no match-qinq-dot1p

Context config>service>vpls>sap>ingress
config>service>ies>if>sap>ingress
config>service>vprn>sub-if>grp-if>sap>ingress

Description This command specifies which dot1Q tag position dot1P bits in a QinQ encapsulated packet should be used to evaluate dot1P QoS classification.

The **match-qinq-dot1p** command allows the top or bottom PBits to be used when evaluating the applied sap-ingress QoS policy's dot1P entries. The **top** and **bottom** keywords specify which position should be evaluated for QinQ encapsulated packets.

By default, the bottom-most service delineating dot1Q tag's dot1P bits are used. Table 109 defines the default behavior for dot1P evaluation when the **match-qinq-dot1p** command is not executed.

Table 109 **Default QinQ and TopQ SAP Dot1P Evaluation**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits

Table 109 **Default QinQ and TopQ SAP Dot1P Evaluation (Continued)**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ/TopQ	TopQ	TopQ PBits
QinQ/TopQ	TopQ BottomQ	TopQ PBits
QinQ/TopQ	TopQ BottomQ	BottomQ PBits

The **no** form of this command restores the default dot1p evaluation behavior for the SAP.

Default no match-qinq-dot1p (no filtering based on p-bits)

(top or bottom must be specified to override the default QinQ dot1p behavior)

Parameters **top** — The top parameter is mutually exclusive to the bottom parameter. When the **top** parameter is specified, the topmost PBits are used (if existing) to match any dot1p dot1p-value entries. Table 110 defines the dot1p evaluation behavior when the **top** parameter is specified.

Table 110 **Top Position QinQ and TopQ SAP Dot1P Evaluation**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	TopQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ/TopQ	TopQ	TopQ PBits
QinQ/TopQ	TopQ BottomQ	TopQ PBits
QinQ/QinQ	TopQ BottomQ	TopQ PBits

bottom — The **bottom** parameter is mutually exclusive to the **top** parameter. When the bottom parameter is specified, the bottom most PBits are used (if existing) to match any dot1p dot1p-value entries. Table 111 defines the dot1p evaluation behavior when the bottom parameter is specified.

Table 111 **Bottom Position QinQ and TopQ SAP Dot1P Evaluation**

Port / SAP Type	Existing Packet Tags	PBits Used for Match
Null	None	None
Null	Dot1P (VLAN-ID 0)	Dot1P PBits
Null	Dot1Q	Dot1Q PBits
Null	TopQ BottomQ	BottomQ PBits
Null	TopQ (No BottomQ)	TopQ PBits
Dot1Q	None (Default SAP)	None
Dot1Q	Dot1P (Default SAP VLAN-ID 0)	Dot1P PBits
Dot1Q	Dot1Q	Dot1Q PBits
QinQ/TopQ	TopQ	TopQ PBits
QinQ/TopQ	TopQ BottomQ	BottomQ PBits
QinQ/QinQ	TopQ BottomQ	BottomQ PBits

qos

Syntax **qos** *policy-id*
no qos

Context config>service>ies>sub-if>grp-if>sap>egress
config>service>vprn>sub-if>grp-if>sap>egress
config>service>vprn>sub-if>grp-if>sap>ingress

Description Associates a Quality of Service (QoS) policy with an egress Service Access Point (SAP) or IP interface.

QoS egress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP or IP interface. If the *policy-id* does not exist, an error are returned.

The **qos** command is used to associate egress QoS policies. The **qos** command only allows egress policies on SAP or IP interface egress. Attempts to associate a QoS policy of the wrong type returns an error.

Only one ingress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type returns an error.

By default, no specific QoS policy is associated with the SAP or IP interface for egress, so the default QoS policy is used.

The normal behavior is for queues to be created per destination.

The **no** form of this command removes the QoS policy association from the SAP or IP interface, and the QoS policy reverts to the default.

Parameters *policy-id* — The egress policy ID to associate with SAP or IP interface on egress. The policy ID must already exist.

Values 1 to 65535

qos

Syntax **qos** *policy-id* [**shared-queuing**]
no qos

Context config>service>ies>sub-if>grp-if>sap>ingress

Description Associates a Quality of Service (QoS) policy with an ingress Service Access Point (SAP) or IP interface.

QoS ingress policies are important for the enforcement of SLA agreements. The policy ID must be defined prior to associating the policy with a SAP or IP interface. If the *policy-id* does not exist, an error is returned.

This **qos** command is used to associate ingress QoS policies. The **qos** command only allows ingress policies to be associated on SAP or IP interface ingress.

Only one ingress and one egress QoS policy can be associated with a SAP or IP interface at one time. Attempts to associate a second QoS policy of a given type returns an error.

By default, no specific QoS policy is associated with the SAP or IP interface for ingress so the default QoS policy is used.

The normal behavior is for queues to be created per destination. Shared and multipoint shared change this behavior creating either unicast or unicast and mcast shared queues.

The **no** form of this command removes the QoS policy association from the SAP or IP interface, and the QoS policy reverts to the default.

Parameters *policy-id* — The ingress policy ID to associate with SAP or IP interface on ingress. The policy ID must already exist.

Values 1 to 65535

shared-queuing — This keyword can only be specified on SAP ingress. Specify the ingress shared queue policy used by a SAP. When the value of this object is null it means that the SAP uses individual ingress QoS queues, instead of the shared ones.

policer-control-policy

Syntax **policer-control-policy** *policy-name*
no policer-control-policy

Context config>service>ies>sub-if>grp-if>sap>egress
config>service>ies>sub-if>grp-if>sap>ingress

Description This command is used to create, delete, or modify policer control policies. The **policer-control-policy** controls the aggregate bandwidth available to a set of child policers. Once created, the policy can be applied to ingress or egress SAPs. The policy can also be applied to the ingress or egress context of a sub-profile.

The **no** form of this command resets the command to the default setting.

Parameters *policy-name* — Specifies the policer control policy name. Each policer control policy name must be unique and adhere to the system policy ASCII naming requirements. If the defined policy name already exists, the system enters that policy's context for editing purposes. If policy-name does not exist, the system attempts to create a policy with the specified name.

qinq-mark-top-only

Syntax [**no**] **qinq-mark-top-only**

Context config>service>ies>sub-if>grp-if>sap>egress
config>service>vprn>if>sap>egress
config>service>vprn>sub-if>grp-if>sap>egress

Description When the encapsulation type is qinq for the access port for the specified SAP, enabling this command specifies which P-bits or DEI bit to mark during packet egress. Only the P-bits or DEI bit in the top Q tag are marked. When this command is disabled, both sets of P-bits and the DEI bit are marked.

scheduler-policy

Syntax **scheduler-policy** *scheduler-policy-name*
no scheduler-policy

Context config>service>ies>sub-if>grp-if>sap>egress
config>service>ies>sub-if>grp-if>sap>ingress
config>service>vprn>sub-if>grp-if>sap>egress

config>service>vprn>sub-if>grp-if>sap>ingress

Description This command applies an existing scheduler policy to an ingress or egress scheduler used by ingress SAP queues or egress SAP policers and queues associated with this multi-service customer site. The schedulers defined in the scheduler policy can only be created once the customer site has been appropriately assigned to a chassis port, channel or slot. Scheduler policies are defined in the **config>qos>scheduler-policy** *scheduler-policy-name* context.

The **no** form of this command removes the configured ingress or egress scheduler policy from the multi-service customer site. When the policy is removed, the schedulers created due to the policy are removed also making them unavailable for the SAP policers or queues associated with the customer site. Policers and queues that lose their parent scheduler association are deemed to be orphaned and are no longer subject to a virtual scheduler. The SAPs that have policers or queues reliant on the removed schedulers enter into an operational state depicting the orphaned status of one or more policers or queues. When the **no** form of this command executed, the customer site ingress or egress node will not contain an applied scheduler policy.

Parameters *scheduler-policy-name:* — Specifies the scheduler policy name to apply to an existing scheduler policy that was created in the **config>qos>scheduler-policy** *scheduler-policy-name* context to create the hierarchy of ingress or egress virtual schedulers. The scheduler names defined within the policy are created and made available to any ingress queues or egress policers and queues created on associated SAPs.

Values Any existing valid scheduler policy name.

host

Syntax **host ip** *ip-address* [**mac** *ieee-address*]] [**subscriber** *sub-ident-string*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*]
no host {[**ip** *ip-address*] [**mac** *ieee-address*]}
no host all

Context config>service>ies>sub-if>grp-if>sap
config>service>ies>if>sap
config>service>ies>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>sap

Description This command creates a static subscriber host for the SAP. Static subscriber hosts may be used by the system for various purposes. Applications within the system that make use of static host entries include anti-spoof filters and ARP cache population.

Multiple static hosts may be defined on the SAP. Each host is identified by either a source IP address, a source MAC address or both a source IP and source MAC address. Every static host definition must have at least one address defined, IP or MAC.

Static hosts can exist on the SAP even with anti-spoof and ARP populate features disabled. When enabled, each feature has different requirements for static hosts.

The **no** form of this command removes the values from the configuration.

Parameters **anti-spoof** — When enabled, this feature uses static and dynamic host information to populate entries into an anti-spoof filter table. The anti-spoof filter entries generated the same type as specified in the **anti-spoof type** parameter. If the SAP anti-spoof filter is defined as **ip**, each static host definition must specify an IP address. If the SAP anti-spoof filter is defined as **ip-mac**, each static host definition must specify both an IP address and MAC address. If definition of a static host is attempted without the appropriate addresses specified for the enabled anti-spoof filter, the static host definition fails.

arp-populate — When enabled, this feature uses static and dynamic host information to populate entries in the system ARP cache.

Attempting to define a static subscriber host that conflicts with an existing DHCP lease state table entry fails.

Use the **no** form of this command to remove a static entry from the system. The specified *ip-address* and *mac-address* must match the host's exact IP and MAC addresses as defined when it was created. When a static host is removed from the SAP, the corresponding anti-spoof entry and/or ARP cache entry is also removed.

ip-address — Specifies this optional parameter when defining a static host. The IP address must be specified for **anti-spoof ip**, **anti-spoof ip-mac** and **arp-populate**. Only one static host may be configured on the SAP with a given IP address.

mac-address — Specifies this optional parameter when defining a static host. The MAC address must be specified for **anti-spoof ip-mac** and **arp-populate**. Multiple static hosts may be configured with the same MAC address given that each definition is distinguished by a unique IP address.

sub-ident-string — Specify this optional parameter to specify an existing subscriber identification profile to be associated with the static subscriber host. The subscriber identification profile is configured in the **config>subscr-mgmt>sub-ident-pol** context. The subscriber information is used by the VPRN SAP arp-reply-agent to determine the proper handling of received ARP requests from subscribers.

For VPRN SAPs with **arp-reply-agent** enabled with the optional *sub-ident* parameter, the static subscriber hosts sub-ident-string is used to determine whether an ARP request received on the SAP is sourced from a host belonging to the same subscriber as the destination host. When both the destination and source hosts from the ARP request are known on the SAP and the subscriber identifications do not match, the ARP request may be forwarded to the rest of the VPRN destinations.

If the static subscriber hosts *sub-ident* string is not defined, the host is not considered to belong to the same subscriber as another host on the SAP.

If source or destination host is unknown, the hosts are not considered to belong to the same subscriber. (ARP messages from unknown hosts are subject to anti-spoof filtering rules applied at the SAP.)

If *sub-ident* is not enabled on the SAP arp-reply-agent, subscriber identification matching is not performed on ARP requests received on the SAP.

ARP requests are never forwarded back to the same SAP or within the receiving SAP's Split Horizon Group.

sub-profile-name — Specify this optional parameter to specify an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

sla-profile-name — Specify this optional parameter to specify an existing SLA profile name to be associated with the static subscriber host. The SLA profile is configured in the **config>subscr-mgmt>sla-profile** context.

ingress

Syntax **ingress**

Context config>service>ies>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>sap

Description This command enables the context to configure ingress SAP Quality of Service (QoS) policies and filter policies.

If no SAP ingress QoS policy is defined, the system default sap-ingress QoS policy is used for ingress processing. If no ingress filter is defined, no filtering is performed.

multi-service-site

Syntax [**no**] **multi-service-site** *customer-site-name*

Context config>service>ies>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>sap

Description This command creates a new customer site or edits an existing customer site with the *customer-site-name* parameter. A customer site is an anchor point to create an ingress and egress virtual scheduler hierarchy. When a site is created, it must be assigned to a chassis slot or port. When scheduler policies are defined for ingress and egress, the scheduler names contained in each policy are created according to the parameters defined in the policy. Multi-service customer sites exist for the sole purpose of creating a virtual scheduler hierarchy and making it available to queues on multiple Service Access Points (SAPs).

The scheduler policy association with the customer site normally prevents the scheduler policy from being deleted until after the scheduler policy is removed from the customer site. The multi-service-site object generates a log message indicating that the association was deleted due to scheduler policy removal.

When the multi-service customer site is created, an ingress and egress scheduler policy association does not exist. This does not prevent the site from being assigned to a chassis slot or prevent service SAP assignment. After the site has been created, the ingress and egress scheduler policy associations can be assigned or removed at anytime.

The **no** form of this command removes the value from the configuration.

Parameters *customer-site-name* — Each customer site must have a unique name within the context of the customer. If *customer-site-name* already exists for the customer ID, the CLI context changes to that site name for the purpose of editing the site scheduler policies or assignment. Any modifications made to an existing site affects all SAPs associated with the site. Changing a scheduler policy association may cause new schedulers to be created and existing policers and queues on the SAPs to no longer be orphaned. Existing schedulers on the site may cease to exist, causing policers and queues relying on that scheduler to be orphaned.

If the *customer-site-name* does not exist, it is assumed that an attempt is being made to create a site of that name in the customer ID context. The success of the command execution depends on the following:

The maximum number of customer sites defined for the chassis slot has not been met.

The *customer-site-name* is valid.

The **create** keyword is included in the command line syntax (if the system requires it).

When the maximum number of customer sites has been exceeded a configuration error occurs, the command will not execute and the CLI context will not change.

If the *customer-site-name* is invalid, a syntax error occurs, the command will not execute and the CLI context will not change.

Values Valid names consist of any string, up to 32 characters, composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

static-host

Syntax **static-host ip** *ip-prefix*[*/prefix-length*] [**mac** *ieee-address*] [**create**]
no static-host ip *ip-prefix*[*/prefix-length*] [**mac** *ieee-address*]
no static-host all [**force**]
no static-host ip *ip-prefix*[*/prefix-length*]

Context config>service>ies>if>sap
config>service>vpls>sap
config>service>ies>sub-if>grp-if>sap
config>service>vprn>if>sap
config>service>vprn>sub-if>grp-if>sap

Description This command creates a static subscriber host for the SAP. Static subscriber hosts may be used by the system for various purposes. Applications within the system that make use of static host entries include anti-spoof, ARP reply agent and source MAC population into the VPLS forwarding database.

Multiple static hosts may be defined on the SAP. Each host is identified by either a source IP address, a source MAC address or both a source IP and source MAC address. Every static host definition must have at least one address defined, IP or MAC.

Static hosts can exist on the SAP even with anti-spoof and ARP reply agent features disabled. When enabled, each feature has different requirements for static hosts.

The **no** form of this command removes a static entry from the system. The specified *ip-address* and *mac-address* must match the host's exact IP and MAC addresses as defined when it was created. When a static host is removed from the SAP, the corresponding anti-spoof filter entry and/or FDB entry is also removed.

Parameters *ip-prefix[/prefix-length* — Specifies information for the specified IP address and mask.

mac-address — Specifies a MAC address. The MAC address must be specified for **anti-spoof mac**, and **anti-spoof ip-mac** and arp-reply-agent (a**rp-reply-agent** is supported by the 7450 ESS only). Multiple static hosts may be configured with the same MAC address given that each definition is distinguished by a unique IP address.

Every static host definition must have at least one address defined, IP or MAC.

force — Specifies the forced removal of the static host addresses.

create — Keyword used to create the static host instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

ancp-string

Syntax **ancp-string** *ancp-string*
no ancp-string

Context config>service>ies>if>sap>static-host
config>service>vpls>sap>static-host
config>service>ies>sub-if>grp-if>sap>static-host
config>service>vprn>if>sap>static-host
config>service>vprn>sub-if>grp-if>sap>static-host

Description This command specifies the ANCP string associated to this SAP host.

The **no** form of this command reverts to the default.

Parameters *ancp-string* — Specifies the ANCP string up to 63 characters.

app-profile

Syntax **app-profile** *app-profile-name*
no app-profile

Context config>service>ies>if>sap>static-host
config>service>vpls>sap>static-host
config>service>ies>sub-if>grp-if>sap>static-host
config>service>vprn>if>sap

config>service>vprn>sub-if>grp-if>sap

Description This command specifies an application profile name.

The **no** form of this command reverts to the default.

Parameters *app-profile-name* — Specifies the application profile name up to 32 characters.

inter-dest-id

Syntax **inter-dest-id** *intermediate-destination-id*
no inter-dest-id

Context config>service>ies>if>sap>static-host
config>service>vpls>sap>static-host
config>service>vprn>if>sap>static-host
config>service>vprn>sub-if>grp-if>sap>static-host

Description This command specifies to which intermediate destination (for example a DSLAM) this host belongs.

The **no** form of this command reverts to the default.

Parameters *intermediate-destination-id* — Specifies the intermediate destination ID.

managed-routes

Syntax **managed-routes**

Context config>service>ies>sub-if>grp-if>sap>static-host
config>service>vprn>sub-if>grp-if>sap>static-host

Description This command enables the context to configure managed route parameters.

route

Syntax **route** {*ip-prefix/length* | *ip-prefix netmask*} [**create**]
no route {*ip-prefix/length* | *ip-prefix netmask*} [**metric** *metric-value*]

Context config>service>ies>sub-if>grp-if>sap>static-host>managed-routes
config>service>vprn>sub-if>grp-if>sap>static-host>managed-routes

Description This command assigns a managed route to a given subscriber-host. As a consequence, a static-route pointing subscriber-host IP address as a next hop is installed in FDB. Up to 16 managed routes per subscriber-host can be configured.

The **no** form of this command removes the respective route.

Parameters *ipv6-prefix/length* | *ipv6-prefix netmask* — Associates an IPv6 managed route to the IPv6 static host. The IPv6 managed routes can overlap with the static host IPv6 address.

ipv4-prefix/length | *ipv6-prefix netmask* — Associates an IPv4 managed route to the IPv4 static host.

Note: A maximum of 16 managed routes can be associated to a static host. IPv4 hosts can only have IPv4 managed routes and IPv6 hosts can only have IPv6 managed routes.

metric-value — A metric can be associated with the provisioned managed route

static-host-mgmt

Syntax **static-host-mgmt**

Context config>service>ies>sub-if>grp-if>sap
config>service>vprn>sub-if>grp-if>sap

Description This command enables the context to configure common parameters for static hosts.

mac-learning-options

Syntax [**no**] mac-learning-options

Context config>service>ies>sub-if>grp-if>sap>static-host-mgmt
config>service>vprn>sub-if>grp-if>sap>static-host-mgmt

Description This command configures additional methods by which the BNG learns the subscriber host MAC.

data-triggered

Syntax [**no**] data-triggered

Context config>service>ies>sub-if>grp-if>sap>static-host-mgmt>mac-learning-options
config>service>vprn>sub-if>grp-if>sap>static-host-mgmt>mac-learning-options

Description This command enables learning of MAC addresses from data packets.

The **no** form of this command disables learning of MAC addresses from data packets.

single-mac

Syntax [**no**] **single-mac**

Context config>service>ies>sub-if>grp-if>sap>static-host-mgmt>mac-learning-options
config>service>vprn>sub-if>grp-if>sap>static-host-mgmt>mac-learning-options

Description This command controls how the SAP learns the IPv6 static host MAC address. Enabling this command indicates that this particular SAP only has one subscriber and only has one MAC address for all hosts. With this parameter enabled, the subscriber's NS and RS source MAC address is used to automatically populate the subscriber MAC address. To allow this auto-populate behavior, the subscriber's NS and RS source IP must be of type link local address.

route

Syntax **route** {*ip-prefix/length* | *ip-prefix netmask*} [**create**]
no route {*ip-prefix/length* | *ip-prefix netmask*} [**metric** *metric-value*]

Context config>service>ies>sub-if>grp-if>sap>static-host>managed-routes
config>service>vprn>sub-if>grp-if>sap>static-host>managed-routes

Description This command assigns a managed route to a given subscriber-host. As a consequence, a static-route pointing subscriber-host IP address as a next hop is installed in FDB. Up to 16 managed routes per subscriber-host can be configured.

The **no** form of this command removes the respective route.

Parameters *ipv6-prefix/length* | *ipv6-prefix netmask* — Associates an IPv6 managed route to the IPv6 static host. The IPv6 managed routes can overlap with the static host IPv6 address.

ipv4-prefix/length | *ipv6-prefix netmask* — Associates an IPv4 managed route to the IPv4 static host.

Note: A maximum of 16 managed routes can be associated to a static host. IPv4 hosts can only have IPv4 managed routes and IPv6 hosts can only have IPv6 managed routes.

metric-value — A metric can be associated with the provisioned managed route

ATM Commands

Note: The commands described in this section apply only to the 7750 SR.

atm

Syntax **atm**

Context config>service>ies>sub-if>grp-if>sap
config>service>vpls>sap
config>service>vprn>if>sap
config>service>vprn>sub-if>grp-if>sap

Description This command the context to configure ATM-related parameters. This command can only be used when a given context (for example, a channel or SAP) supports ATM functionality such as:

- Configuring ATM port or ATM port-related functionality on MDAs supporting ATM functionality
- Configuring ATM-related configuration for ATM-based SAPs that exist on MDAs supporting ATM functionality.

If ATM functionality is not supported for a given context, the command returns an error.

egress

Syntax **egress**

Context config>service>ies>sub-if>grp-if>sap>atm
config>service>vpls>sap>atm
config>service>vprn>if>sap>atm
config>service>vprn>sub-if>grp-if>sap>atm

Description This command enables the context to configure egress ATM attributes for the SAP.

encapsulation

Syntax **encapsulation** *atm-encap-type*

Context config>service>ies>sub-if>grp-if>sap>atm
config>service>vprn>if>sap>atm
config>service>vprn>sub-if>grp-if>sap>atm

Description This command configures RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*, encapsulation for an ATM PVCC delimited SAP.

This command specifies the data encapsulation for an ATM PVCC delimited SAP. The definition references RFC 2684 and to the ATM Forum LAN Emulation specification.

Ingress traffic that does not match the configured encapsulation are dropped.

The **no** form of this command removes the value from the configuration.

Default The encapsulation is driven by the services for which the SAP is configured. For IES service SAPs, the default is aal5snap-routed.

Parameters *atm-encap-type* — Specify the encapsulation type.

Values **aal5snap-routed** — Routed encapsulation for LLC encapsulated circuit (LLC/SNAP precedes protocol datagram) as defined in RFC 2684.

aal5mux-ip — Routed IP encapsulation for VC multiplexed circuit as defined in RFC 2684

ingress

Syntax **ingress**

Context config>service>ies>sub-if>grp-if>sap>atm
config>service>vprn>if>sap>atm
config>service>vprn>sub-if>grp-if>sap>atm

Description This command configures ingress ATM attributes for the SAP.

traffic-desc

Syntax **traffic-desc** *traffic-desc-profile-id*
no traffic-desc

Context config>service>ies>sub-if>grp-if>sap>atm>egress
config>service>ies>sub-if>grp-if>sap>atm>ingress
config>service>vpls>sap>atm>ingress
config>service>vpls>sap>atm>egress
config>service>vprn>if>sap>atm>egress
config>service>vprn>if>sap>atm>ingress
config>service>vprn>sub-if>grp-if>sap>atm>egress
config>service>vprn>sub-if>grp-if>sap>atm>ingress

Description This command assigns an ATM traffic descriptor profile to a given context (for example, a SAP). When configured under the ingress context, the specified traffic descriptor profile defines the traffic contract in the forward direction. When configured under the egress context, the specified traffic descriptor profile defines the traffic contract in the backward direction.

The **no** form of this command reverts the traffic descriptor to the default traffic descriptor profile.

Default The default traffic descriptor (trafficDescProfileId. = 1) is associated with newly created PVCC-delimited SAPs.

Parameters *traffic-desc-profile-id* — Specify a defined traffic descriptor profile (see the QoS atm-td-profile command).

oam

Syntax **oam**

Context config>service>ies>sub-if>grp-if>sap>atm
config>service>vpls>sap>atm
config>service>vprn>interface >sap>atm
config>service>vprn>sub-if>grp-if>sap>atm

Description This command enables the context to configure OAM functionality for a PVCC delimiting a SAP.

The ATM-capable MDAs support F5 end-to-end OAM functionality (AIS, RDI, Loopback):

- ITU-T Recommendation I.610 - B-ISDN Operation and Maintenance Principles and Functions version 11/95
- GR-1248-CORE - Generic Requirements for Operations of ATM Network Elements (NEs). Issue 3 June 1996
- GR-1113-CORE - Bellcore, Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1, July 1994

alarm-cells

Syntax [**no**] **alarm-cells**

Context config>service>ies>sub-if>grp-if>sap>atm>oam
config>service>vpls>sap>atm
config>service>vprn>if>sap>atm>oam
config>service>vprn>sub-if>grp-if>sap>atm>oam

Description This command configures AIS/RDI fault management on a PVCC. Fault management allows PVCC termination to monitor and report the status of their connection by propagating fault information through the network and by driving PVCCs operational status.

When alarm-cells functionality is enabled, PVCCs operational status is affected when a PVCC goes into AIS or RDI state because of an AIS/RDI processing (that is assuming nothing else affects PVCCs operational status, PVCC goes DOWN, when it enters a fault state and comes back UP, when it exits that fault state) and RDI cell are generated when PVCC is operationally DOWN. No OAM-specific SNMP trap is raised whenever an endpoint enters/exits an AIS or RDI states, however, if as result of an OAM state change, the PVCC changes operational status, then a trap is expected from an entity the PVCC is associated with (for example a SAP).

The **no** form of this command disables alarm-cells functionality for a PVCC. When alarm-cells functionality is disabled, PVCCs operational status is no longer affected by PVCCs OAM state changes due to AIS/RDI processing (when alarm-cells is disabled, a PVCC changes operational status to UP, if it was DOWN because of the alarm-cell processing) and RDI cells are not generated as result of PVCC going into AIS or RDI state, however, PVCCs OAM status records OAM faults as described above.

Default Enabled for PVCCs delimiting IES SAPs

periodic-loopback

Syntax [**no**] **periodic-loopback**

Context config>service>ies>sub-if>grp-if>sap>atm>oam
config>service>vprn>if >sap>atm>oam
config>service>vprn>sub-if>grp-if>sap>atm

Description This command enables periodic OAM loopbacks on this SAP. This command is only configurable on IES and VPRN SAPs. When enabled, an ATM OAM loopback cell is transmitted every period as configured in the **config>system>atm>oam>loopback-period** *period* context.

If a response is not received and consecutive retry-down retries also result in failure, the endpoint transitions to an alarm indication signal/loss of clock state. Then, an ATM OAM loopback cell is transmitted every period as configured in the loopback-period *period*. If a response is received for the periodic loopback and consecutive retry-up retries also each receive a response, the endpoint transitions back to the up state.

The **no** form of this command sets the value back to the default.

9.25.2.1.27 Vport Commands

ethernet

Syntax **ethernet**

Context config>port

Description This command the context to configure Ethernet port attributes.

This context can only be used when configuring Fast Ethernet, gigabit or 10-G Fast Ethernet or Ethernet LAN ports on an appropriate MDA.

egress-scheduler-override

Syntax [**no**] **egress-scheduler-override**

Context config>port>ethernet

Description This command applies egress scheduler overrides. When a port scheduler is associated with an egress port, it is possible to override the following parameters:

- The **max-rate** allowed for the scheduler.
- The maximum **rate** for each priority level 8 through 1.
- The CIR associated with each priority level 8 through 1.

See the SR OS Quality of Service Guide for command syntax and usage for the **port-scheduler-policy** command.

The **no** form of this command removes all override parameters from the egress port or channel scheduler context. Once removed, the port scheduler reverts all rate parameters back to the parameters defined on the **port-scheduler-policy** associated with the port.

level

Syntax **level** *priority-level* **rate** *pir-rate* [**cir** *cir-rate*]
no level *priority-level*

Context config>port>ethernet>egress-scheduler-override

Description This command overrides the maximum and CIR rate parameters for a specific priority level on the port or channel's port scheduler instance. When the **level** command is executed for a priority level, the corresponding priority **level** command in the **port-scheduler-policy** associated with the port is ignored. The override **level** command supports the keyword **max** for the **rate** and **cir** parameter. When executing the **level** override command, at least the **rate** or **cir** keywords and associated parameters must be specified for the command to succeed.

The **no** form of this command removes the local port priority level rate overrides. Once removed, the port priority level uses the port scheduler policies level command for that priority level.

Parameters *priority-level* — Identifies which of the eight port priority levels are being overridden.

Values 1 to 8

pir-rate — Overrides the port scheduler policy's maximum level rate and requires either the **max** keyword or a rate defined in kilobits-per-second to follow.

Values 1 to 40000000, max

cir-rate — Overrides the port scheduler policy's within-cir level rate and requires either the **max** keyword or a rate defined in kilobits-per-second to follow.

Values 0 to 40000000, max

max — Removes any existing rate limit imposed by the port scheduler policy for the priority level allowing it to use as much total bandwidth as possible.

access

Syntax **access**

Context config>port>ethernet

Description This command configures Ethernet access port parameters.

egress

Syntax **egress**

Context config>port>ethernet>access

Description This command enables the context to configure Ethernet access egress port parameters.

vport

Syntax **vport** *name* [**create**]
no vport *name*

Context config>port>ethernet>access>egress

Description This command configures a scheduling node, referred to as virtual port, within the context of an egress Ethernet port. The Vport scheduler operates either like a port scheduler with the difference that multiple Vport objects can be configured on the egress context of an Ethernet port, or it can be an aggregate rate when an egress port-scheduler policy is applied to the port.

The Vport is always configured at the port level even when a port is a member of a LAG.

When a port scheduler policy is applied to a Vport the following command is used:

configure>port>ethernet>access>egress>vport>port-scheduler-policy *port-scheduler-policy-name*

The CLI does not allow the user to apply a port scheduler policy to a Vport if one has been applied to the port. Conversely, the CLI does not allow the user to apply a port scheduler policy to the egress of an Ethernet port if one has been applied to any Vport defined on the access egress context of this port. The **agg-rate-limit**, along with an egress **port-scheduler**, can be used to ensure that a given Vport does not oversubscribe the port's rate.

SAP and subscriber host queues can be port-parented to a Vport scheduler in a similar way they port-parent to a port scheduler or can be port-parented directly to the egress port-scheduler if the agg-rate-limit is used.

When the Vport uses an aggregate rate, the following command is used:

configure>port>ethernet>access>egress>vport>agg-rate-limit

The **no** form of this command removes the Vport name from the configuration.

Parameters *name* — Specifies the name of the Vport scheduling node and can be up to 32 ASCII characters. This does not need to be unique within the system but is unique within the port or a LAG.

create — Keyword required to create the configuration context.

agg-rate-limit

Syntax **agg-rate-limit** *agg-rate*
no agg-rate-limit

Context config>port>ethernet>access>egress>vport

Description This command configures an aggregate rate for the Vport. This command is mutually exclusive with the **port-scheduler-policy** command.

The **no** form of this command reverts to the default.

Parameters *agg-rate* — Specifies the rate limit for the Vport.

Values max, 1 to 10000000

egress-rate-modify

Syntax [**no**] **egress-rate-modify**

Context config>port>ethernet>access>egress>vport

Description This command is used to apply HQoS Adjustment to a Vport. HQoS Adjustment refers to the dynamic adjustment of the rate limit at an QoS enforcement point within router when the multicast traffic stream is disjointed from the unicast traffic stream. This QoS enforcement point within router represents the physical point further down in the access part of the network where the two streams join each other and potentially can cause congestion.

An example would be a PON port which is shared amongst subscriber's multicast traffic (single copy of each channel) and subscriber's unicast traffic. The bandwidth control point for this PON port resides in the upstream router BNG node in the form of a Vport. In case that the multicast delivery method in the router BNG utilizes redirection, the multicast traffic in the router BNG flows outside of the subscriber or the Vport context and thus bypasses any bandwidth enforcement in the router. To correct this, a Vport bandwidth adjustment is necessary in the router that accounts for the multicast bandwidth consumption that is bypassing Vport in the router but is present in the PON port whose bandwidth is controlled by Vport.

An estimate of the multicast bandwidth consumption on the PON port can be made at the Vport level based on the IGMP messages sourced from the subscribers behind the PON port. This process is called HQoS Adjustment.

A multicast channel bandwidth is subtracted from or added to the Vport rate limit according to the received IGMP Join/Leave messages and the channel bandwidth definition policy associated with the Vport (indirectly through a group-interface). Since the multicast traffic on the PON port is shared amongst subscribers behind this PON port, only the first IGMP Join or the last IGMP Leave per multicast channel is tracked for the purpose of the Vport bandwidth modification.

The Vport rate that is affected by this functionality depends on the configuration:

- In case the **agg-rate-limit** within the Vport is configured, its value is modified based on the IGMP activity associated with the subscriber under this Vport.
- In case the **port-scheduler-policy** within the Vport is referenced, the **max-rate** defined in the corresponding **port-scheduler-policy** is modified based on the IGMP activity associated with the subscriber under this Vport.

The channel bandwidth definition policy is defined in the mcac policy in the **configure>router>mcac>policy** context. The policy is applied under the group interface or in case of redirection under the redirected interface.

The rates in effect can be displayed with the following two commands:

show port 1/1/5 vport *name*

qos scheduler-hierarchy port *port-id* vport *vport-name*

The configuration of a scheduler policy under a Vport, which is only applicable to Ethernet interfaces, is mutually exclusive with the configuration of the **egress-rate-modify** parameter.

The configuration of a scheduler policy under a Vport, which is only applicable to Ethernet interfaces, is mutually exclusive with the configuration of the **egress-rate-modify** parameter.

The **no** form of this command reverts to the default.

host-match

Syntax **host-match dest** *destination-string* [**create**]
no host-match dest *destination-string*

Context config>port>ethernet>access>egr>qgrp

Description This command configures host matching for the Ethernet port egress queue-group.

The **no** form of this command removes the destination string from the configuration.

Parameters *destination-string* — Specify a host match destination string up to 32 characters.

create — Keyword used to create the host match. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

port-scheduler-policy

Syntax **port-scheduler-policy** *port-scheduler-policy-name*
no port-scheduler-policy

Context config>port>ethernet>access>egress>vport

Description This command specifies the destination and organization strings to be used for matching subscriber hosts with this Vport.

The parent Vport of a subscriber host queue, which has the port-parent option enabled, is determined by matching the destination string dest string associated with the subscriber and the organization string org string associated with the subscriber host with the strings defined under a Vport on the port associated with the subscriber.

If a given subscriber host policers or queue does not have the port-parent option enabled, it is foster-parented to the Vport used by this subscriber and which is based on matching the dest string and org string. If the subscriber could not be matched with a Vport on the egress port, the host policer or queue will not be bandwidth controlled and competes for bandwidth directly based on its own PIR and CIR parameters.

By default, a subscriber host policer or queue with the port-parent option enabled is scheduled within the context of the port's port scheduler policy.

The **agg-rate rate**, **port-scheduler-policy** and **scheduler-policy** commands are mutually exclusive. Changing between the use of a scheduler policy and the use of an **agg-rate** or **port-scheduler-policy** involves removing the existing command and applying the new command. Applying a scheduler policy to a Vport is only applicable to Ethernet interfaces.

The **no** form of this command removes the *port-scheduler-policy-name* from the configuration.

The **agg-rate** *rate*, **port-scheduler-policy** and **scheduler-policy** commands are mutually exclusive. Changing between the use of a scheduler policy and the use of an **agg-rate**/**port-scheduler-policy** involves removing the existing command and applying the new command.

The **no** form of this command reverts to the default.

Parameters *port-scheduler-policy-name* — Specifies an existing **port-scheduler-policy** configured in the **config>qos** context.

scheduler-policy

Syntax **scheduler-policy** *scheduler-policy-name*
no scheduler-policy

Context config>port>ethernet>access>egress>vport

Description This command specifies a scheduler policy to associate to the Vport. Scheduler policies are configured in the **configure>qos>scheduler>policy** context. Each scheduler policy is divided up into groups of schedulers based on the tier each scheduler is created under. A tier is used to give structure to the schedulers within a policy and define rules for parent scheduler associations. The policy defines the hierarchy and operating parameters for virtual schedulers.

The **no** form of this command removes the configured egress scheduler policy from the Vport.

The **agg-rate rate**, **port-scheduler-policy** and **scheduler-policy** commands are mutually exclusive. Changing between the use of a scheduler policy and the use of an agg-rate/port-scheduler-policy involves removing the existing command and applying the new command.

The configuration of a scheduler policy under a Vport is mutually exclusive with the configuration of the egress-rate-modify parameter.

The **no** form of this command reverts to the default.

Parameters *scheduler-policy-name* — Specifies the *scheduler-policy-name* parameter applies an existing scheduler policy that was created in the **config>qos>scheduler-policy** *scheduler-policy-name* context to create the hierarchy of egress virtual schedulers.

parent-location

Syntax **parent-location** {**default** | **sla**}
no parent-location

Context config>qos>sap-egress

Description This command determines the expected location of the parent schedulers for queues configured with a parent command within the SAP egress policy. All parent schedulers must be configured within a scheduler policy applied at the location corresponding to the parent-location parameter.

If a parent scheduler name does not exist at the specified location, the queue is not parented and is orphaned.

The **no** form of this command reverts to the default.

Default parent-location default

Parameters **default** — Specifies that the queues should look for their in the SAP's egress scheduler-policy or the multi-service-site's egress scheduler policy depending on which is configured and Subscriber queues looks for their parents in the subscriber-profile's egress scheduler-policy.

When the SAP egress policy is applied to an SLA profile for a subscriber, the parent schedulers of the queues must be configured in the scheduler policy applied to the subscriber's SUB profile.
When the SAP egress policy is applied to a SAP, the parent schedulers of the queues need to be configured in the scheduler policy applied to the SAP or the multi-service site.

sla — Specifies that the queues should look for their parent in the SLA-profile's egress scheduler-policy; all queues that are not port-parented and to which SLA-profile is not applicable behaves as orphans in this case.

When the SAP egress policy is applies to an SLA profile for a subscriber, the parent schedulers of the queues must be configured in the scheduler policy applied to the same SLA profile.
If this parameter is configured within a SAP egress policy that is applied to any object except of the egress of an SLA profile, the configured parent schedulers are not found and so the queues are not parented and are orphaned.

parent-location

Syntax **parent-location** {**none** | **sub** | **vport**}
no parent-location

Context config>qos>scheduler-policy

Description This command determines the expected location of the parent schedulers for the tier 1 schedulers configured with a parent command within the scheduler policy. The parent schedulers must be configured within a scheduler policy applied at the location corresponding to the parent location parameter.

If a parent scheduler name does not exist at the specified location, the schedulers are not parented and are orphaned.

The configuration of **parent-location** and **frame-based-accounting** commands in a scheduler policy is mutually exclusive in to ensure consistency between the different scheduling levels.

The **no** form of this command reverts to the default.

Default parent-location none

Parameters **none** — Indicates that the tier 1 schedulers do not have a parent scheduler and the configuration of the parent under a tier 1 scheduler is blocked. Conversely, this parameter is blocked when any tier 1 scheduler has a parent configured.

sub — When the scheduler policy is applied to an SLA profile for a subscriber, the parent schedulers of the tier 1 schedulers need to be configured in the scheduler policy applied to the subscriber's SUB profile.

If this parameter is configured within a scheduler policy that is applied to any object except for the egress of an SLA profile, the configured parent schedulers are found and so the tier 1 schedulers are not parented and are orphaned.

vport — When the scheduler policy is applied to an SLA profile, a SUB profile for a subscriber or to the egress of a PW SAP, the parent schedulers of the tier 1 schedulers need to be configured in the scheduler policy applied to the Vport to which the subscriber is assigned.

If this parameter is configured within a scheduler policy that is applied to any object except for the egress of an SLA profile or SUB profile, or to the egress of a PW SAP, the configured parent schedulers is not found and so the tier 1 schedulers will not be parented and is orphaned. This parameter is not supported when policers-hqos-manageable is configured in the SAP egress QoS policy.

9.25.2.1.28 IPoE Session Commands

ipoe-session-policy

Syntax **ipoe-session-policy** *policy-name* [**create**]
no ipoe-session-polic*y policy-name*

Context config>subscr-mgmt

Description This command configures an IPoE session policy. The policies are referenced from subscriber interfaces, group interfaces and capture SAPs. Multiple IPoE session policies can be configured.

The **no** form of this command removes the policy name from the configuration.

Parameters *policy-name* — Specifies the IPoE policy name up to 32 characters.

create — Keyword required to create the configuration context

circuit-id-from-auth

Syntax [**no**] **circuit-id-from-auth**

Context config>subscr-mgmt>ipoe-policy

Description This command takes the circuit ID value from the authentication server to identify the session.

The **no** form of this command reverts to the default.

description

Syntax **description** *description-string*
no description

Context config>subscr-mgmt>ipoe-policy

Description This command takes the *circuit-id* value from the authentication server when it creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes any description string from the context.

Parameters *description-string* — Specifies a text string describing the entity. Allowed values are any string up to 80 characters, composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, and so on), the entire the entire string must be enclosed within double quotes.

session-key

Syntax **session-key sap mac** [**cid**] [**rid**]
no session-key

Context config>subscr-mgmt>ipoe-policy

Description This command configures the key to logically group subscriber hosts that belong to the same dual stack end device in an IPoE session.

The SAP and MAC address are always part of the IPoE session key. Optionally the Circuit-Id/Interface-Id or Remote-Id can be added to the session key.

The **no** form of this command reverts to the default.

Default sap mac

Parameters **sap** — Includes the SAP as part of the IPoE session key. The **sap** parameter is mandatory and cannot be removed from the key.

mac — Includes the MAC address as part of the IPoE session key. The **mac** parameter is mandatory and cannot be removed from the key.

cid — Optionally adds the DHCPv4 Relay Agent Circuit-Id (Option 82, sub Option 1) and DHCPv6 Interface-Id (Option 18) field to the IPoE session key.

rid — Optionally adds the DHCPv4 Relay Agent Remote-Id (Option 82, sub Option 2) and DHCPv6 Remote-Id (Option 37) field to the IPoE session key. For DHCPv6, the enterprise number is excluded from the key.

sap and **mac** are mandatory parameters while **cid** and **rid** are optional and mutually exclusive. Valid IPoE session key parameters are: **sap mac**, **sap mac cid** and **sap mac rid**.

session-timeout

Syntax **session-timeout** *timeout*
no session-timeout

Context config>subscr-mgmt>ipoe-policy

Description This command defines the time in seconds between 1 second and 360 days before the IPoE session is disconnected. The default value is unlimited session timeout.

The **no** form of this command reverts to the default.

Parameters *timeout* — Specifies the session timeout in seconds.

Values 1 to 31104000

sap

Syntax **sap** *sap-id* [**create**] [**capture-sap**]

Context config>service>vpls

Description This command declares a given SAP as a primary (or secondary) VPLS port.

The **no** form of this command reverts to the default.

Parameters **capture-sap** — Specifies a capturing SAP in which triggering packets is sent to the CPM. Non-triggering packets captured by the capture SAP is dropped. This parameter applies to the 7450 ESS or 7750 SR only.

create — Specifies to create a SAP instance. The create keyword requirement can be enabled or disabled in the **environment>create** context.

ipoe-session

Syntax [**no**] **ipoe-session**

Context config>service>vpls>sap
config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if
config>service>ies>sub-if
config>service>vprn>sub-if

Description This command configures IPoE session parameters.

force-auth

Syntax **force-auth** [**cid-change**] [**rid-change**]
force-auth disabled
no force-auth

Context config>service>ies>sub-if>grp-if>ipoe-session
config>service>vprn>sub-if>grp-if>ipoe-session

Description By default, if the circuit-id/interface-id or remote-id in the IPoE session re-authentication trigger packet (such as a DHCP renewal) is not empty and different from the circuit-id/interface-id or remote-id stored in the IPoE session data, a forced re-authentication is performed, ignoring the configured **min-auth-interval**. This default behavior can be changed with this command.

The **no** form of this command reverts to the default behavior.

Default force-auth cid-change rid-change
force-auth disabled on wlan-gw group interfaces

Parameters **cid-change** — Perform a forced re-authentication upon a circuit-id/interface-id change. An empty circuit-id/interface-id is not considered a change.

rid-change — Perform a forced re-authentication upon a remote-id change. an empty remote-id is not considered a change. For DHCPv6, the enterprise number is excluded from the comparison.

disabled — Specifies that the **min-auth-interval never** is ignored. The system does not perform a forced re-authentication upon a circuit-id or interface-id or remote-id change.

ipoe-session-policy

Syntax **ipoe-session-policy** *policy-name*
no ipoe-session-policy

Context config>service>vpls>sap> ipoe-session
config>service>ies>sub-if>grp-if>ipoe-session
config>service>vprn>sub-if>grp-if>ipoe-session

Description This command specifies the IPoE session policy applicable for this group interface or capture SAP.

On WLAN GW group interfaces, it is not possible to change this value.

The **no** form of this command reverts to the default.

Default no ipoe-session-policy
ipoe-session-policy default on WLAN GW group interfaces

Parameters *policy-name* — Specifies the IPoE session policy name up to 32 characters.

min-auth-interval

Syntax **min-auth-interval** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
min-auth-interval infinite
no min-auth-interval

Context config>service>ies>sub-if>grp-if>ipoe-session
config>service>vprn>sub-if>grp-if>ipoe-session

Description Re-authentication for IPoE sessions enable dynamic policy changes.

This command configures the maximum frequency of re-authentications by specifying a minimum interval between two non-forced authentications for the same IPoE session.

A forced authentication is by default triggered by a circuit-id, interface-id or remote-id change (see the force-auth command).

Re-authentications are, by default, disabled and can be enabled by configuring a **min-auth-interval**.

Setting the **min-auth-interval** to zero seconds always re-authenticates on each trigger packet.

The **no** form of this command reverts to the default behavior.

Default min-auth-interval infinite

Parameters *days* — Specifies the min number of days between two non-forced authentications for IPoE sessions.

Values 0 to 365

hours — Specifies the min number of hours between two non-forced authentications for IPoE sessions.

Values 0 to 23

minutes — Specifies the min number of minutes between two non-forced authentications for IPoE sessions.

Values 0 to 59

seconds — Specifies the min number of seconds between two non-forced authentications for IPoE sessions.

Values 0 to 59

infinite — Specifies that non-forced re-authentications for IPoE sessions is not performed.

radius-session-timeout

Syntax **radius-session-timeout** {**backwards-compatible** | **ignore** | **absolute**}
no radius-session-timeout

Context config>service>ies>sub-if>grp-if>ipoe-session
config>service>vprn>sub-if>grp-if>ipoe-session

Description This command specifies how to interpret the session-timeout coming from a RADIUS VSA in an Access-Accept or CoA message.

The value of this command can only be changed on wlan-gw group interfaces.

The **no** form of this command to resets the default behavior.

Default radius-session-timeout absolute (backward compatible on wlan-gw group interfaces)

Parameters **backwards-compatible** — Specifies that the VSA is interpreted as an IPv4 lease time if the Alc-Lease-Time attribute is not present and an absolute timeout otherwise. The VSA is treated the same as for non-ipoe session DHCP hosts.

ignore — Specifies that the VSA meaning is irrelevant for IPoE session and should be ignored.

absolute — Specifies that the VSA would be treated as a timeout starting from the moment the IPoE session is set up.

sap-session-limit

Syntax **sap-session-limit** *sap-session-limit*
no sap-session-limit

Context config>service>ies>sub-if>grp-if>ipoe-session
config>service>vprn>sub-if>grp-if>ipoe-session

Description This command specifies the number of IPoE sessions per SAP allowed for this group-interface

The **no** form of this command reverts to the default.

Default sap-session-limit 1

Parameters *sap-session-limit* — Specifies the number of allowed IPoE sessions.

Note: The operational maximum value may be smaller due to equipped hardware dependencies.

Values 1 to 131071
131071 on wlan-gw group interfaces

session-limit

Syntax **session-limit** *session-limit*
no session-limit

Context config>service>ies>sub-if>grp-if>ipoe-session
config>service>vprn>sub-if>grp-if>ipoe-session
config>service>ies>sub-if>ipoe-session
config>service>vprn>sub-if>ipoe-session

Description This command specifies the number of IPoE sessions allowed for this group interface or retail subscriber interface.

The **no** form of this command reverts to the default.

Default session-limit 1

Parameters *session-limit* — Specifies the number of allowed IPoE sessions.

Note: The operational maximum value may be smaller due to equipped hardware dependencies.

Values 1 to 131071
1 to 500000 (retail subscriber interface)

user-db

Syntax **user-db** *local-user-db-name*
no user-db

Context config>service>vpls>sap>ipoe-session

config>service>ies>sub-if>grp-if>ipoe-session
config>service>vprn>sub-if>grp-if>ipoe-session

Description This command configures the local user database to use for IPoE session authentication.

When configured on a capture SAP, the group interface must have the same local user database configured.

On a wlan-gw group interface, the **no** form of this command indicates that the user database is picked from the following sources in the order shown:

1. dhcp
2. ipv6>dhcp6
3. ipv6>router-solicit

If no user database can be found in any of these locations, processing continues as if no user database was configured. This behavior is for backwards compatibility reasons only; when using a LUDB, it should be explicitly added to the IPoE session configuration.

The **no** form of this command reverts to the default.

Parameters *local-user-db-name* — Specifies the local user database name up to 32 characters

shutdown

Syntax [**no**] **shutdown**

Context config>service>vpls>sap> ipoe-session
config>service>ies>sub-if>grp-if>ipoe-session
config>service>vprn>sub-if>grp-if>ipoe-session

Description The **shutdown** command enables or disables IPoE session management on a group interface or capture SAP.

A **shutdown** of the IPoE session CLI hierarchy on a group-interface clears all active IPoE sessions on that interface, resulting in a deletion of all corresponding subscriber hosts.

On wlan-gw group interfaces it is not possible to disable an IPoE session.

The **no** form of this command reverts to the default.

Default shutdown
no shutdown on wlan-gw group interfaces

9.25.2.1.29 SHCV Policy Commands

shcv-policy

Syntax **shcv-policy** *name* [**create**]
no shcv-policy *name*

Context config>subscr-mgmt

Description This command configures a Subscriber Host Connectivity Verification (SHCV) policy. An SHCV policy can be applied to both the subscriber management group interface and VPLS instances. All SHCV-related features inside a group interface and a VPLS service follows the configuration specified in the SHCV policy. The SHCV policy and the SHCV configuration on a group interface are mutually exclusive. Only one can be applied to the group interface.

The **no** form of this command removes the policy name from the configuration.

Parameters *name* — Specifies the name of the SHCV policy, up to 32 characters.

create — Keyword required to create the configuration context.

layer-3

Syntax **layer-3**

Context config>subscr-mgmt>shcv-policy

Description This command enables the context to configure SHCV behavior parameters for IES and VPRN services.

source-ip-origin

Syntax **source-ip-origin** {**interface** | **vrrp**}
no source-ip-origin

Context config>subscr-mgmt>shcv-policy>layer-3

Description This command selects the source IP address to be used for SHCV messages.

The **no** form of this command reverts to the default.

Parameters **interface** — Specifies to use the interface IP as the source address of SHCV.

vrrp — Specifies to use the VRRP configured IP as the source address of SHCV.

periodic

Syntax **periodic**

Context config>subscr-mgmt>shcv-policy

Description This command enables the context to configure periodic SHCV properties for the subscriber management group-interface. This tool periodically scans all known DHCP hosts only and perform unicast ARP/NS requests. The subscriber host connectivity verification maintains state (connected versus not-connected) for all hosts.

action

Syntax **action** {**alarm** | **remove**}
no action

Context config>subscr-mgmt>shcv-policy>periodic

Description This command configures the action to take when the periodic connectivity verification failed.

The **no** form of this command reverts to the default.

Default action alarm

Parameters **alarm** — Raises an alarm indicating that the host is disconnected.

remove — Raises an alarm and releases all allocated resources (addresses, prefixes, queues, table entries, and so on). Static hosts are removed.

interval

Syntax **interval** *minutes*
no interval

Context config>subscr-mgmt>shcv-policy>periodic

Description This command specifies the time interval which all known sources should be verified. The actual rate is dependent on the number of known hosts and intervals.

The **no** form of this command reverts to the default.

Default interval 30 minutes

Parameters *minutes* — Specifies the interval, in minutes, between periodic connectivity checks.

Values 1 to 6000

retry-count

Syntax **retry-count** [*count*]
no retry-count

Context config>subscr-mgmt>shcv-policy>periodic

Description This command configures the number of retransmissions.

The **no** form of this command reverts to the default.

Default retry-count 1 — For trigger-type ip-conflict, host-limit-exceeded and mobility

retry-count 10 — For trigger-type inactivity and MAC learning

Parameters *count* — Specifies the number of retransmissions in periodic connectivity verification.

Values 2 to 29

timeout

Syntax **timeout** *seconds*
no timeout

Context config>subscr-mgmt>shcv-policy>periodic

Description This command configures the timeout before a retransmission in triggered connectivity verification.

The **no** form of this command reverts to the default.

Default timeout 10

Parameters *seconds* — Specifies the timeout, in seconds, before a retransmission in triggered connectivity verification.

Values 10 to 60

trigger

Syntax **trigger** *trigger-type*

Context config>subscr-mgmt>shcv-policy

Description This command enables to context to configure SHCV triggers.

Parameters *trigger-type* — Specifies the trigger SHCV properties for the subscriber management group-interface.

Values

ip-conflict — Upon detecting an IP conflict for the new host, a trigger SHCV is sent using a unicast ARP/NS. The request verifies the old host's connectivity to the BNG. An unresponsive host is removed from the system, allowing the new host to connect. The new host must resend its address request in order to be created as an ESM host in the system.

host-limit-exceeded — Upon exceeding the host-limit for (**sla-profile**, **ipoe-session** *sap-session-limit*, **ipoe-session** *session-limit*, and **arp-host**), a trigger SHCV is sent. The request verifies the old host's connectivity to the BNG. An unresponsive host is removed from the system, allowing the new host to connect. The new host must resend its address request in order to be created as an ESM host in the system.

inactivity — A trigger SHCV is sent to the idle host to verify the host's connectivity to the BNG. An unresponsive host is removed from the system.

mobility — Detects detects an IP or MAC conflict between different SAPs. Upon detecting a MAC or an IP conflict, a trigger SHCV is sent. The request verifies the old host's connectivity to the BNG. An unresponsive host is removed from the system, allowing the new host to connect. The new host must resend its address request in order to be created as an ESM host in the system.

mac-learning — Specifies the trigger SHCV properties to learn the MAC address of static hosts. Upon a **no shutdown**, a trigger SHCV is sent to the host to learn the host's MAC address.

retry-count

Syntax **retry-count** [*count*]
no retry-count

Context config>subscr-mgmt>shcv-policy>trigger

Description This command configures the number of retransmissions in periodic connectivity verification.

The **no** form of this command reverts to the default.

Default retry-count 1

Parameters *count* — Specifies the number of retransmissions in periodic connectivity verification.

Values 1 to 29

timeout

Syntax **timeout** *seconds*
no timeout

Context config>subscr-mgmt>shcv-policy>trigger

Description This command configures the timeout before a retransmission.

The **no** form of this command reverts to the default.

Default timeout 1 — trigger-type ip-conflict, host-limit-exceeded and mobility

timeout 2 — trigger-type inactivity and mac-learning

Parameters *seconds* — Specifies the retry timeout in seconds.

Values 1 to 60

vpls

Syntax **vpls** *service-id*

Context config>subscr-mgmt>shcv-policy

Description This command enables the context to configure SHCV behavior in VPLS services. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* for VPLS service command syntax and descriptions.

vprn

Syntax **vprn** *service-id*

Context config>service

Description This command enables the context to configure SHCV behavior in VPRN services. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* for VPRN service command syntax and descriptions.

ies

Syntax **ies** *service-id*

Context config>service

Description This command enables the context to configure SHCV behavior in IES services. Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* for IES service command syntax and descriptions.

shcv-policy-ipv4

Syntax **shcv-policy-ipv4** *name*
no shcv-policy-ipv4

Context config>service>vpls
config>service>vpls>sap
config>service>ies>if
config>service>vprn>if
config>service>vprn>sub-if>grp-if
config>service>ies>sub-if>grp-if

Description This command references the shcv-policy to be used exclusive for IPv4 subscriber hosts. The shcv-policy name must already exist in the **config>subscr-mgmt** context.

The **no** form of this command removes the policy name from the service configuration.

Parameters *name* — Specifies an existing SHCV policy name up to 32 characters.

shcv-policy-ipv6

Syntax **shcv-policy-ipv6** *name*
no shcv-policy-ipv6

Context config>service>vprn>sub-if>grp-if
config>service>ies>sub-if>grp-if

Description This command references the SHCV policy to be used exclusive for IPv6 subscriber hosts. The policy name must already exist in the **config>subscr-mgmt** context.

The **no** form of this command removes the policy name from the service configuration.

Parameters *name* — Specifies an existing SHCV policy name up to 32 characters.

shcv-policy

Syntax **shcv-policy** *name*
no shcv-policy

Context config>service>vprn>sub-if>grp-if
config>service>ies>sub-if>grp-if

Description This command references the SHCV policy to be used for both IPv4 and IPv6 subscriber hosts. The policy name must already exist in the **config>subscr-mgmt** context.

The **no** form of this command removes the policy name from the service configuration.

Parameters *name* — Specifies an existing SHCV policy name up to 32 characters.

source-ip

Syntax **source-ip** *ip-address*
no source-ip

Context config>subscr-mgmt>shcv-policy>vpls

Description This command configures the IPv4 address to be used as source address for connectivity verification in a VPLS service.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the IPv4 address to be used as source address for connectivity verification in a VPLS service.

source-mac

Syntax **source-mac** *ieee-address*
no source-mac

Context config>subscr-mgmt>shcv-policy>vpls

Description Specifies the MAC address to be used as source address for connectivity verification in a VPLS service.

The **no** form of this command reverts to the default.

Parameters *ieee-address* — Specifies the 48-bit MAC address in the xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx format.

9.25.2.1.30 Service DHCP and Anti-Spoof Filtering Commands

mac-linking

Syntax **mac-linking** *ip-address*
no mac-linking

Context config>service>ies>sub-if>grp-if>sap>static-host
config>service>vprn>sub-if>grp-if>sap>static-host

Description This command associates this IPv6 host to the specified IPv4 host through the learned MAC address. A learned MAC from the IPv6 host is associated with the IPv4 host and vice versa.

The **no** form of this command removes the IP address from the configuration.

rip-policy

Syntax **rip-policy** *rip-policy-name*
no rip-policy

Context config>service>ies>sub-if>grp-if>sap>static-host

Description This command specifies the name of the RIP policy up to 32 characters.

The **no** form of this command removes the policy name from the static-host configuration.

igmp-host-tracking

Syntax **igmp-host-tracking**

Context config>service>vprn>sub-if>grp-if>sap

Description This command enables the context to configure IGMP host tracking parameters.

disable-router-alert-check

Syntax [**no**] **disable-router-alert-check**

Context config>service>vprn>sub-if>grp-if>sap>igmp-host-tracking

Description This command disables the IGMP router alert check option.

The **no** form of this command enables the router alert check.

expiry-time

Syntax **expiry-time** *expiry-time*
no expiry-time

Context config>service>vprn>sub-if>grp-if>sap>igmp-snooping

Description This command configures the time that the system continues to track inactive hosts.

The **no** form of this command removes the values from the configuration.

Parameters *expiry-time* — Specifies the time, in seconds, that this system continues to track an inactive host.

Values 1 to 65535

import

Syntax **import** *policy-name*
no import

Context config>service>vprn>sub-if>grp-if>sap>igmp-snooping

Description This command specifies the import routing policy to be used for IGMP packets to be used on this SAP or SDP. Only a single policy can be imported on a single SAP at any time.

The **no** form of this command removes the policy association from the SAP or SDP.

Parameters *policy-name* — Specifies the routing policy name. Allowed values are any string up to 32 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes. Routing policies are configured in the config>router>policy-options context The router policy must be defined before it can be imported.

max-num-group

Syntax **max-num-groups** *max-num-groups*
no max-num-groups

Context config>service>vprn>sub-if>grp-if>sap>igmp-snooping

Description This command configures the maximum number of multicast groups allowed to be tracked.

The **no** form of this command removes the values from the configuration.

Parameters *max-num-groups* — Specifies the maximum number of multicast groups allowed to be tracked.

Values 1 to 196607

max-num-sources

Syntax **max-num-sources** *max-num-sources*
no max-num-sources

Context config>service>vprn>sub-if>grp-if>sap>igmp-snooping

Description This command configures the maximum number of multicast sources allowed to be tracked per group.

The **no** form of this command removes the value from the configuration.

Parameters *max-num-sources* — Specifies the maximum number of multicast sources allowed to be tracked per group.

Values 1 to 1000

max-num-grp-sources

Syntax [**no**] **max-num-grp-sources** [*number*]

Context config>service>vprn>sub-if>grp-if>sap>igmp-snooping

Description This command configures the maximum number of group sources for which IGMP can have local receiver information based on received IGMP reports on this interface. When this configuration is changed dynamically to a value lower than currently accepted number of group sources, the group sources that are already accepted are not deleted. Only new group sources are not allowed. When this object has a value of 0, there is no limit to the number of group sources.

The **no** form of this command removes the value from the configuration.

Parameters *number* — Specifies the maximum number of multicast sources allowed to be tracked per group.

Values 1 to 32000

9.25.2.1.31 Auto-Generated Subscriber Identification Key Commands

auto-sub-id-key

Syntax **auto-sub-id-key**

Context config>subscr-mgmt

Description This command enables the context to configure auto-generated subscriber identification key parameters.

ipoe-sub-id-key

Syntax **ipoe-sub-id-key** *sub-id-key* [*sub-id-key*]
no ipoe-sub-id-key

Context config>subscr-mgmt>auto-sub-id-key

Description This command enables certain fields to become the base for auto-generation of the default sub-id name. The sub-id name is auto generated if there is not a more specific method available. Such more specific methods would be a default sub-id name as a sap-id, a preconfigured static string or explicit mappings based on RADIUS/LUDB returned strings.

In case that a more specific sub-id name generation method is not available and the auto-id keyword is defined under the def-sub-id hierarchy, the sub-id name is generated by concatenating fields defined in this command separated by a "|" character.

The maximum sub-id name length is 32 characters while the concatenation of subscriber identification fields can easily exceed 32 characters. Subscriber host instantiation will fail in case that the sub-id name is based on subscriber identification fields whose concatenated length exceeds 32 characters. Failing the host creation rather than truncating sub-id name on a 32 character boundary will prevent collision of sub-ids (subscriber name duplication).

In case that a more specific sub-id name generation method is not available and the auto-id keyword is not defined under the def-sub-id hierarchy, the sub-id name is a random 10 character encoded string based on the fields defined under this command.

There is only one set of identification fields allowed per host type (IPoE or PPP) per chassis.

The **no** form of this command reverts to the default.

Default ipoe-sub-id-key mac sap-id

Parameters *sub-id-key* — Specifies the auto-generated sub-id keys for IPoE hosts.

Values **mac** — Specifies that the MAC address can be combined with other subscriber host identification fields (circuit-id, remote-id, session-id or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the mac address is used as a concatenation field in the sub-id name, then its format becomes a string xx:xx:xx:xx:xx:xx with the length 17B.

The MAC address as the subscriber host identification field is not applicable to PPPoA hosts or static hosts.

circuit-id — Specifies that the circuit-id can be combined with other subscriber host identification fields (mac, remote-id, session-id or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the circuit-id is used as a concatenation field in the sub-id name, then its format becomes access-node-id eth slot/port:[vlan-id] or access-node-id atm slot/port:vpi.vci with a variable length.

Note: If circuit-id contains any non-printable ASCII characters, the entire circuit-id string is formatted in hex in the sub-id name output. Otherwise all characters in circuit-id is converted to ASCII. ASCII printable characters contain bytes in range 0x20 to 0x7E.

The circuit-id as the subscriber identification field is not applicable to PPPoA hosts, ARP hosts or static hosts.

remote-id — Specifies that the remote-id can be combined with other subscriber host identification fields (mac, circuit-id, session-id or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the remote-id is used as a concatenation field in the sub-id name, then its format becomes a remote-id string with a variable length.

Note: If remote-id contains any non printable ASCII characters, the entire remote-id string is formatted in hex in the sub-id name output. Otherwise all characters in remote-id is converted to ASCII. ASCII printable characters contain bytes in range 0x20 to 0x7E.

The remote-id as the subscriber identification field is not applicable to PPPoA hosts, ARP hosts or static hosts.

sap-id — Specifies that the sap-id can be combined with other subscriber host identification fields (mac, circuit-id, remote-id, or session-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

If a circuit-id is used as a concatenation field in the sub-id name, then its format becomes: slot/mda:[outer-vlan].[inner-vlan] with a variable length.

The sap-id as the subscriber identification field is applicable to all hosts types with exception of static hosts.

dual-stack-remote-id — Specifies that the dual stack remote ID is used as the base for the auto-generated subscriber identification

ppp-sub-id-key

Syntax **ppp-sub-id-key** *sub-id-key* [*sub-id-key*]
no ppp-sub-id-key

Context config>subscr-mgmt>auto-sub-id-key

Description This command enables certain fields to become the base for auto-generation of default sub-id name. The sub-id name is auto-generated if there is not a more specific method available. Examples of these specific methods would be a default sub-id name as a sap-id, a preconfigured static string or explicit mappings based on RADIUS/LUDB returned strings.

In case that a more specific sub-id name generation method is not available and the **auto-id** keyword is defined under the def-sub-id hierarchy, the sub-id name is generated by concatenating fields defined in this command separated by a “|” character.

The maximum sub-id name length is 32 characters while the concatenation of subscriber identification fields can easily exceed 32 characters. The subscriber host instantiation will fail if the sub-id name is based on subscriber identification fields whose concatenated length exceeds 32 characters. Failing the host creation rather than truncating sub-id name on a 32 character boundary will prevent collision of sub-ids (subscriber name duplication).

In case that a more specific sub-id name generation method is not available and the **auto-id** keyword is not defined under the def-sub-id hierarchy, the sub-id name is a random 10 character encoded string based on the fields defined under this command.

There is only one set of identification fields allowed per host type (IPoE or PPP) per chassis.

The **no** form of this command reverts to the default.

Default ppp-sub-id-key mac sap-id session-id

Parameters *sub-id-key* — Specifies the auto-generated sub-id keys for PPP hosts.

Values **mac** — Specifies that the MAC address can be combined with other subscriber host identification fields (circuit-id, remote-id, session-id or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the mac address is used as a concatenation field in the sub-id name, then its format becomes a string xx:xx:xx:xx:xx:xx with the length 17B.

The MAC address as the subscriber host identification field is not applicable to PPPoA hosts or static hosts.

circuit-id — Specifies that the circuit-id can be combined with other subscriber host identification fields (mac, remote-id, session-id or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

If the circuit-id is used as a concatenation field in the sub-id name, then its format becomes access-node-id eth slot/port:[vlan-id] or access-node-id atm slot/port:vpi.vci with a variable length.

Note: If circuit-id contains any non printable ASCII characters, the entire circuit-id string is formatted in hex in the sub-id name output. Otherwise all characters in circuit-id is converted to ASCII. ASCII printable characters contain bytes in range 0x20 to 0x7E.

If the circuit-id is used as the subscriber identification field is not applicable to PPPoA hosts, ARP hosts or static hosts.

remote-id — Specifies that the remote-id can be combined with other subscriber host identification fields (mac, circuit-id, session-id or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

If the remote-id is used as a concatenation field in the sub-id name, then its format becomes a remote-id string with a variable length.

Note: If the remote ID contains any non printable ASCII characters, the entire remote-id string is formatted in hex in the sub-id name output. Otherwise all characters in remote-id is converted to ASCII. ASCII printable characters contain bytes in range 0x20 to 0x7E.

The remote ID as the subscriber identification field is not applicable to PPPoA hosts, ARP hosts or static hosts.

sap-id — The SAP ID can be combined with other subscriber host identification fields (mac, circuit-id, remote-id, or session-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the circuit-id is used as a concatenation field in the sub-id name, then its format becomes: slot/mda:[outer-vlan].[inner-vlan] with a variable length.

The SAP ID as the subscriber identification field is applicable to all hosts types with exception of static hosts.

session-id — The session ID can be combined with other subscriber host identification fields (mac, circuit-id, remote-id, or sap-id) to form a sub-id name in a user readable format or as a random 10 character encoded value.

In case that the circuit ID is used as a concatenation field in the sub-id name, then its format becomes a decimal number with variable length.

The session ID as the subscriber identification field is applicable only to PPPoE/PPPoEoA type hosts.

9.25.2.1.32 IPv6 Commands

ipv6

Syntax **ipv6**

Context config>service>ies>if
config>service>ies>sub-if
config>service>ies>sub-if>grp-if
config>service>vprn>if
config>service>vprn>sub-if
config>service>vprn>sub-if>grp-if
config>service>ies>sub-if>lcl-addr-assign
config>service>vprn>sub-if>lcl-addr-assign

Description This command enables the context to configure IPv6 parameters.

address

Syntax **address** *ip-address/mask* [*netmask*] [**gw-ip-address** *gw-ip-address*] [**populate-host-routes**] [**track-srrp** *srrp-instance*] [**holdup-time** *milli-seconds*]

no address *ip-address/mask* [*netmask*]

Context config>service>vprn>sub-if
config>service>vprn>sub-if

Description This command configures the subscriber interface address along with additional parameters related to multi-chassis redundancy.

The **no** form of this command reverts to the default.

Parameters *ip-address* — The IP address of the IP interface. The *ip-address* portion of the **address** command specifies the IP host address that is used by the IP interface within the subnet. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

/ — The forward slash is a parameter delimiter and separates the *ip-address* portion of the IP address from the mask that defines the scope of the local subnet. No spaces are allowed between the *ip-address*, the "*/*" and the *mask-length* parameter. If a forward slash is not immediately following the *ip-address*, a dotted decimal mask must follow the prefix.

mask — The subnet mask in dotted decimal notation. When the IP prefix is not specified in CIDR notation, a space separates the *ip-address* from a traditional dotted decimal mask. The *mask* parameter indicates the complete mask that is used in a logical AND function to derive the local subnet of the IP address. Allowed values are dotted decimal addresses in the range 128.0.0.0 – 255.255.255.252.

Note: A mask of 255.255.255.255 is reserved for system IP addresses.

netmask — The subnet mask in dotted decimal notation.

Values 0.0.0.0 - 255.255.255.255

gw-ip-address — Specifies a separate IP address within the subnet for SRRP routing purposes. This parameter must be followed by a valid IP interface that exists within the subscriber subnet created by the address command. The defined gateway IP address cannot currently exist as a subscriber host (static or dynamic). If the defined ip-address already exists as a subscriber host address, the address command will fail. The specified ip-address must be unique within the system.

The *gw-ip-address* parameter may be specified at anytime. If the subscriber subnet was created previously, executing the **address** command with a *gw-ip-address* parameter will simply add the SRRP gateway IP address to the existing subnet.

If the **address** command is executed without the *gw-ip-address* parameter when the subscriber subnet is associated with an active SRRP instance, the address will fail. If the SRRP instance is inactive or removed, executing the **address** command without the *gw-ip-address* parameter removes the SRRP gateway IP address from the specified subscriber subnet.

If the **address** command is executed with a new GW address, all SRRP instances associated with the specified subscriber subnet is updated with the new SRRP gateway IP address.

populate-host-routes — Specifies to populate subscriber-host routes in local FDB. Storing them in FDB benefits topologies only where the external router advertises more specific routes than the one corresponding to locally configured subscriber-interface subnets.

milli-seconds — Specifies the time to wait, in milli-seconds, for the route before it accepts the new state attribute. This timer is used to prevent fluctuations in route advertisement caused by short lived SRRP instabilities, in the case that such condition arises.

Values 100 to 5000

srrp-inst — Enables the subscriber interface route to track the SRRP state of the specified SRRP instance. The route updates its state attribute to reflect the *state* of SRRP instance:

- Master = srrp-master
- Any other = srrp-non-master

Routing policy can be applied towards the state attribute in order to customize the advertisement of the route. Only one SRRP instance can be tracked per subscriber interface route. Tracked SRRP instance can be part of the Fate Sharing Group. This command can be enabled at any time.

allow-multiple-wan-addresses

Syntax [**no**] **allow-multiple-wan-addresses**

Context config>service>ies>sub-if>ipv6
config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6
config>service>vprn>sub-if>grp-if>ipv6

Description This command enables host to have two WAN addresses, one from DHCP IA_NA and one from SLAAC assignment.

The **no** form of this command reverts to the default.

allow-unmatching-subnets

Syntax [**no**] **allow-unmatching-subnets**

Context config>service>ies>sub-if>ipv6
config>service>vprn>sub-if>ipv6

Description This command allows address assignment for IPoEv6 and PPPoEv6 hosts in cases where the subscriber host assigned IPv6 address or prefix falls outside of the subscriber-prefix range explicitly configured for the subscriber-interface (**configure>service>vprn/ies>sub-if>ipv6**) or the subscriber-prefix is not configured at all.

SLAAC hosts are installed in the FDB as /64 entries, the length of the installed DHCP-PD prefix is dictated by the prefix-length and the DHCP-NA host is installed as /128 entries.

IPv4 subscriber hosts are unaffected by this command.

The **no** form of this command reverts to the default.

bfd

Syntax **bfd** *transmit-interval* [**receive** *receive-interval*] [**multiplier** *multiplier* [**echo-receive** *echo-interval*] [**type cpm-np**]
no bfd

Context config>service>ies>if>ipv6
config>service>ies>sub-if>ipv6
config>service>vprn>if>ipv6

Description This command specifies the BFD parameters for the associated IP interface. If no parameters are defined the default value are used.

The multiplier specifies the number of consecutive BFD messages that must be missed from the peer before the BFD session state is changed to down and the upper level protocols (OSPF, IS-IS, BGP or PIM) is notified of the fault.

The **no** form of this command removes BFD from the interface.

Note: On the 7750 SR, the *transmit-interval,* **receive** *receive-interval, and* **echo-receive** *echo-interval* values can only be modified to a value less than 100 when:

1. The **type cpm-np option** is explicitly configured.
2. The service is shut down (**shutdown**)
3. The interval is specified 10 to 100000.
4. The service is re-enabled (**no shutdown**)

To remove the **type cpm-np** option, re-issue the **bfd** command without specifying the **type** parameter.

Parameters *transmit-interval* — Sets the transmit interval for the BFD session.

Values 100 to 100000

10 to 100000 (for the 7750 SR only; see the note above)

Default 100

receive *receive-interval* — Sets the receive interval for the BFD session.

Values 100 to 100000
10 to 100000 (for the 7750 SR only; see the note above)

Default 100

multiplier *multiplier* — Sets the multiplier for the BFD session.

Values 3 to 20

Default 3

echo-receive *echo-interval* — Sets the minimum echo receive interval, in milliseconds, for the BFD session.

Values 100 to 100000
10 to 100000 (for the 7750 SR only; see the Note above)

Default 100

type cpm-np — For the 7750 SR only, specifies that BFD sessions associated with this interface is created on the CPM network processor to allow for fast timers down to 10 ms granularity.

dad-disable

Syntax [**no**] **dad-disable**

Context config>service>ies>if>ipv6
config>service>vprn>if>ipv6

Description This command disables duplicate address detection (DAD) on a per-interface basis. This prevents the router from performing a DAD check on the interface. All IPv6 addresses of an interface with DAD disabled, immediately enter a preferred state, without checking for uniqueness on the interface. This is useful for interfaces which enter a looped state during troubleshooting and operationally disable themselves when the loop is detected, requiring manual intervention to clear the DAD violation.

The **no** form of this command turns off **dad-disable** on the interface.

default-dns

Syntax **default-dns** *ip-address* [**secondary** *ip-address*]
no default-dns

Context config>service>ies>sub-if
config>service>vprn>sub-if

Description This command configures last resort IP DNS addresses that can be used for name resolution by IPoE hosts (IA_NA, IA_PD and SLAAC) and PPPoE hosts (IA_NA, IA_PD and SLAAC).

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the IP address of the primary DNS server.

secondary *ip-address* — Specifies the IP address of the secondary DNS server (optional).

default-dns

Syntax **default-dns** *ipv6-address* [**secondary** *ipv6-address*]
no default-dns

Context config>service>ies>sub-if>ipv6
config>service>vprn>sub-if>ipv6

Description This command configures last resort IPv6 DNS addresses that can be used for name resolution by IPoEv6 hosts (IA_NA, IA_PD and SLAAC) and PPPoEv6 hosts (IA_NA, IA_PD and SLAAC).

The **no** form of this command reverts to the default.

Parameters *ipv6-address* — Specifies the IPv6 address of the primary DNS server.

secondary *ipv6-address* — Specifies the IPv6 address of the secondary DNS server (optional).

delegated-prefix-length

Syntax **delegated-prefix-length** *bits*
delegated-prefix-length variable
no delegated-prefix-length

Context config>service>ies>sub-if>ipv6
config>service>vprn>sub-if>ipv6

Description This command configures the subscriber interface level setting for delegated prefix length. The delegated prefix length for a subscriber- interface can be either set to a fixed value that is explicitly configured under the subscriber interface CLI hierarchy or a variable value that can be obtained from various sources. This command can be changed only when no IPv6 prefixes are configured under the **subscriber-interface** context.

The **no** form of this command reverts to the default.

Default no delegated-prefix-length (the delegated prefix length is 64)

Parameters *bits* — Specifies the delegated prefix length in bits. This value is applicable to the entire subscriber interface. In case that the delegated prefix length is also supplied via other means (LUDB, RADIUS or DHCP Server), such supplied value must match the value configured under the **subscriber-interface** context. Otherwise, the prefix instantiation in the router fails.

Values 48 to 64

variable — Specifies that the delegated prefix value can be of any length between 48 to 64. The value itself can vary between the prefixes and is provided at the time of prefix instantiation. The order of priority for the source of the delegated prefix length is:

- LUDB
- RADIUS
- DHCPv6 server

dhcp6-relay

Syntax [**no**] **dhcp6-relay**

Context config>service>ies>if>ipv6

Description This command enables the context to configure DHCPv6 relay parameters for the IES interface.

The **no** form of this command disables DHCP6 relay.

dhcp6-relay

Syntax [**no**] **dhcp6-relay**

Context config>service>ies>if>ipv6
config>service>vprn>if>ipv6

Description This command enables the context to configure DHCPv6 relay parameters for the IES interface.

The **no** form of this command disables DHCPv6 relay.

dhcp6-server

Syntax [**no**] **dhcp6-server**

Context config>service>ies>if>ipv6
config>service>vprn>if>ipv6

Description This command enables the context to configure DHCPv6 server parameters for the IES interface.

The **no** form of this command disables the DHCPv6 server.

icmp6

Syntax **icmp6**

Context config>service>ies>if>ipv6
config>service>vprn>if>ipv6

Description This command configures ICMPv6 parameters for the IES interface.

ipoe-bridged-mode

Syntax [**no**] **ipoe-bridged-mode**

Context config>service>ies>sub-if>ipv6
config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6
config>service>vprn>sub-if>grp-if>ipv6

Description This command enables IPv6 IPoE bridged mode.

The **no** form of this command disables the IPv6 IPoE bridged mode.

packet-too-big

Syntax **packet-too-big** [*number seconds*]
no packet-too-big

Context config>service>ies>if>ipv6>icmp6

Description This command specifies whether packet-too-big ICMP messages should be sent. When enabled, ICMPv6 packet-too-big messages are generated by this interface.

The **no** form of this command disables the sending of ICMPv6 packet-too-big messages.

Default packet-too-big 100 10

Parameters *number* — Specifies the number of ICMP messages that are too large to send in the time frame specified by the *seconds* parameter.

Values 10 to 1000

Default 100

seconds — Specifies the time frame, in seconds, that is used to limit the number of "packet-too-big" ICMP messages issued.

Values 1 to 60

Default 10

param-problem

Syntax **param-problem** [*number seconds*]
no param-problem

Context config>service>ies>if>ipv6>icmp6
config>service>ies>sub-if>grp-if>icmp
config>service>vprn>sub-if>grp-if>icmp

Description This command specifies whether parameter-problem ICMPv6 messages should be sent. When enabled, parameter-problem ICMPv6 messages are generated by this interface.

The **no** form of this command disables the sending of parameter-problem ICMPv6 messages.

Default param-problem 100 10

Parameters *number* — Specifies the number of parameter-problem ICMPv6 messages to send in the time frame specified by the *seconds* parameter.

Values 10 to 1000

Default 100

seconds — Specifies the time frame, in seconds, that is used to limit the number of parameter-problem ICMPv6 messages issued.

Values 1 to 60

Default 10

redirects

Syntax **redirects** [*number seconds*]
no redirects

Context config>service>ies>if>icmp
config>service>ies>if>ipv6>icmp6
config>service>vprn>if
config>service>ies>sub-if>grp-if>icmp

Description This command configures ICMPv6 redirect messages. When enabled, ICMPv6 redirects are generated when routes are not optimal on this router and another router on the same subnetwork has a better route in order to alert that node that a better route is available.

When disabled, ICMPv6 redirects are not generated.

The **no** form of this command reverts to the default.

Default redirects 100 10

Parameters *number* — Specifies the number of version 6 redirects are to be issued in the time frame specified by the *seconds* parameter.

Values 10 to 1000

Default 100

seconds — Specifies the time frame, in seconds, that is used to limit the number of version 6 redirects issued.

Values 1 to 60

Default 10

time-exceeded

Syntax **time-exceeded** [*number seconds*]
no time-exceeded

Context config>service>ies>if>ipv6>icmp6

Description This command specifies whether time-exceeded ICMP messages should be sent. When enabled, ICMPv6 time-exceeded messages are generated by this interface.

When disabled, ICMPv6 time-exceeded messages are not sent.

The **no** form of this command reverts to the default.

Default time-exceeded 100 10

Parameters *number* — Specifies the number of time-exceeded ICMP messages are to be issued in the time frame specified by the *seconds* parameter.

Values 10 to 2000

seconds — Specifies the time frame, in seconds, that is used to limit the number of time-exceeded ICMP message to be issued.

Values 1 to 60

unreachables

Syntax **unreachables** [*number seconds*]
no unreachables

Context config>service>ies>if>ipv6>icmp6

Description This command specifies that ICMPv6 host and network unreachable messages are generated by this interface.

When disabled, ICMPv6 host and network unreachable messages are not sent.

The **no** form of this command reverts to the default.

Default unreachables 100 10

Parameters *number* — Specifies the number of destination unreachable ICMPv6 messages are issued in the time frame specified by the *seconds* parameter.

Values 10 to 2000

seconds — Specifies the time frame, in seconds, that is used to limit the number of destination unreachable ICMPv6 messages to be issued.

Values 1 to 60

local-proxy-nd

Syntax [**no**] **local-proxy-nd**

Context config>service>ies>if>ipv6

Description When this command is enabled, the interface replies to neighbor solicitation requests when both the hosts are on the same subnet. In this case, ICMP redirects are disabled. When this command is disabled, the interface does not reply to neighbor solicitation requests if both the hosts are on the same subnet.

The **no** form of this command reverts to the default.

local-dhcp-server

Syntax **local-dhcp-server** *local-server-name*
no local-dhcp-server

Context config>service>ies>if
config>service>vprn>if>ipv6

Description This command assigns a DHCP server to the interface.

Parameters *local-server-name* — Specifies an existing local server name.

local-proxy-nd

Syntax [**no**] **local-proxy-nd**

Context config>service>ies>if>ipv6
config>service>vprn>if>ipv6

Description This command enables local proxy neighbor discovery on the interface.

The **no** form of this command disables local proxy neighbor discovery.

neighbor

Syntax **neighbor** *ipv6-address mac-address*
no neighbor *ipv6-address*

Context config>service>ies>if>ipv6

Description This command configures IPv6-to-MAC address mapping on the IES interface.

Parameters *ipv6-address* — The IPv6 address of the interface for which to display information.

Values

ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0..FFFF]H
	d - [0..255]D

mac-address — Specifies the 48-bit MAC address for the IPv6-to-MAC address mapping in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

neighbor-limit

Syntax **neighbor-limit** *limit* [**log-only**] [**threshold** *percent*]
no neighbor-limit

Context config>service>ies>if>ipv6

Description This command configures the maximum amount of dynamic IPv6 neighbor entries that can be learned on an IP interface.

When the number of dynamic neighbor entries reaches the configured percentage of this limit, an SNMP trap is sent. When the limit is exceeded, no new entries are learned until an entry expires and traffic to these destinations is dropped. Entries that have already been learned is refreshed.

The **no** form of this command removes the **neighbor-limit**.

Default neighbor-limit 90 percent

Parameters **log-only** — Enables the warning message to be sent at the specified threshold percentage, and also when the limit is exceeded. However, entries above the limit is learned.

percent — The threshold value (as a percentage) that triggers a warning message to be sent.

Values 0 to 100

limit — The number of entries that can be learned on an IP interface expressed as a decimal integer. If the limit is set to 0, dynamic neighbor learning is disabled and no dynamic neighbor entries are learned.

Values 0 to 102400

proxy-nd-policy

Syntax **proxy-nd-policy** *policy-name* [*policy-name*]
no proxy-nd-policy

Context config>service>ies>if>ipv6

Description This command applies a proxy neighbor discovery policy for the interface.

Parameters *policy-name* — Specifies an existing neighbor discovery policy name. Allowed values are any string, up to 32 characters, composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes. The specified policy name must already be defined. Up to 5 policy-names can be specified in a single statement.

qos-route-lookup

Syntax **qos-route-lookup** [**source** | **destination**]
no qos-route-lookup

Context config>service>ies>if
config>service>ies>if>ipv6
config>service>ies>sub-if>grp-if
config>service>ies>sub-if>grp-if>ipv6

Description This command enables QoS classification of the ingress IP packets on an interface based on the QoS information associated with routes in the forwarding table.

If the optional **destination** parameter is specified and the destination address of an incoming IP packet matches a route with QoS information the packet is classified to the forwarding class (fc) and priority associated with that route, overriding the fc and priority/profile determined from the sap-ingress or network qos policy associated with the IP interface. If the destination address of the incoming packet matches a route with no QoS information the fc and priority of the packet remain as determined by the sap-ingress or network qos policy.

If the optional **source** parameter is specified and the source address of an incoming IP packet matches a route with QoS information the packet is classified to the fc and priority associated with that route, overriding the fc and priority/profile determined from the sap-ingress or network qos policy associated with the IP interface. If the source address of the incoming packet matches a route with no QoS information the fc and priority of the packet remain as determined by the sap-ingress or network qos policy.

If neither the optional **source** or **destination** parameter is present, then the default is **destination** address matching.

The functionality enabled by the qos-route-lookup command can be applied to IPv4 packets or IPv6 packets on an interface, depending on whether it is present at the interface context (applies to IPv4) or the interface>ipv6 context (applies to IPv6). Subscriber management group interfaces also do not support the source QPPB option.

The **no** form of this command reverts to the default.

Default destination

Parameters **source** — Enables QoS classification of incoming IP packets based on the source address matching a route with QoS information.

destination — Enables QoS classification of incoming IP packets based on the destination address matching a route with QoS information.

reachable-time

Syntax **reachable-time** *milliseconds*
no reachable-time

Context config>service>ies>sub-if>grp-if>ipv6>router-ad

Description This command configures how long this router should be considered reachable by other nodes on the link after receiving a reachability confirmation.

Parameters *milliseconds* — The length of time the router should be considered reachable for default router selection.

Values 0 to 3600000

router-advertisements

Syntax [**no**] **router-advertisements**

Context config>service>ies>sub-if>ipv6
config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6
config>service>vprn>sub-if>grp-if>ipv6

Description This command enables IPv6 router advertisements for this interface.

The **no** form of this command disables the router advertisements.

router-solicit

Syntax **router-solicit**

Context config>service>ies>sub-if>grp-if>ipv6
config>service>vprn>sub-if>grp-if>ipv6
config>service>vprn>sub-if>ipv6
config>service>ies>sub-if>ipv6

Description This command enables the context to configure parameters used for router-solicit based authentication.

subscriber-prefixes

Syntax **subscriber-prefixes**

Context config>services>ies>sub-if>ipv6
config>services>vprn>sub-if>ipv6

Description This command enables the context to configure aggregate off-link subscriber prefixes associated with this subscriber interface. Individual prefixes are specified under the prefix context list aggregate routes in which the nexthop is indirect via the subscriber interface.

secure-nd

Syntax [**no**] **secure-nd**

Context config>service>ies>if>ipv6

Description This command enables Secure Neighbor Discovery (SeND) on the IPv6 interface.

The **no** form of this command reverts to the default and disabled SeND.

stale-time

Syntax **stale-time** *seconds*
no stale-time

Context config>service>ies>ipv6
config>service>ies>if>ipv6

Description This command configures the time a neighbor discovery cache entry can remain stale before being removed.

The **no** form of this command removes the stale-time value.

Parameters *seconds* — The allowed stale time (in seconds) before a neighbor discovery cache entry is removed.

Values 60 to 65535

tcp-mss

Syntax **tcp-mss** *mss-value*
no tcp-mss

Context config>service>ies>if
config>service>ies>if>ipv6

Description This command statically sets the TCP maximum segment size (MSS) for TCP connections originated from the associated IP interface to the specified value.

The **no** form of this command removes the static value and allows the TCP MSS value to be calculated based on the IP MTU value by subtracting the base IP and TCP header lengths from the IP MTU value (tcp_mss = ip_mtu – 40).

Parameters *mss-value* — The TCP MSS value that should be used in the TCP SYN packet during the three-way handshake negotiation of a TCP connection.

Note: 9158 = max-IP_MTU (9198)-40

Values 536 to 9158 (IPv4)
1220 to 9138 (IPv6)

urpf-check

Syntax [**no**] **urpf-check**

Context config>service>ies>if
config>service>ies>if>ipv6
config>service>vprn>if
config>service>vprn>if>ipv6

Description This command enables the unicast RPF (uRPF) Check on this interface.

The **no** form of this command disables the uRPF Check on this interface.

ignore-default

Syntax [**no**] **ignore-default**

Context config>service>ies>if
config>service>ies>if>ipv6
config>service>ies>sub-if>grp-if>ipv6
config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command enables the default route when performing a uRPF check.

The **no** form of this command disables the default route.

mode

Syntax **mode** {**strict** | **loose** | **strict-no-ecmp**}
no mode

Context config>service>ies>if>urpf-check
config>service>ies>if>ipv6>urpf-check
config>service>vprn>if>urpf-check
config>service>vprn>if>ipv6>urpf-check

Description This command specifies the mode of unicast RPF check.

The **no** form of this command reverts to the default (strict) mode.

Default mode strict

Parameters **strict** — When specified, uRPF checks whether incoming packet has a source address that matches a prefix in the routing table, and whether the interface expects to receive a packet with this source address prefix.

loose — In **loose** mode, uRPF checks whether incoming packet has source address with a corresponding prefix in the routing table. However, the loose mode does not check whether the interface expects to receive a packet with a specific source address prefix. This object is valid only when **urpf-check** is enabled.

strict-no-ecmp — When a packet is received on an interface in this mode and the SA matches an ECMP route the packet is dropped by uRPF.

9.25.2.2 Show Commands

Note: The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

diameter-node

Syntax **diameter-node** *origin-host-string* **peer** *dest-host-string*
diameter-node *origin-host-string* **peer** *dest-host-string* **statistics**
diameter-node *origin-host-string* **peers**
diameter-node *origin-host-string* **routing-table**

Context show>aaa

Description This command displays information about a Diameter node.

Output The following are example of Diameter node information.

Sample Output

```
show aaa diameter node "node1.myrealm.com" routing-table
===============================================================================
Diameter Routing Table (Node: node1.myrealm.com)
===============================================================================
Realm Name                      Appl      Pref:Index    Server-Identifier
-------------------------------------------------------------------------------
subrealm.telenet.be             Gx        30:1          node1.subrealm.telenet.be
telenet.be                      Gx        60            node2.subrealm.telenet.be
telenet.be                      relay     50            node3.subrealm.telenet.be
-------------------------------------------------------------------------------
No. of Routes: 5
===============================================================================
```

Table 112 describes routing entries in the Diameter realm-based routing table fields.

Table 112 **Diameter Node Routing Table Field Descriptions**

Field	Description
Realm Name	The destination realm name
Appl	The application for which this route will be used. Supported applications: • NASREQ • Gx • Gy • Relay (encompass all applications)
Pref:Index	The preference and index of the peer. Peers for the same realm are selected in order of preference. Peers with higher preference are preferred. If multiple peers have the same preference for the same realm, the peer with a lowest index is selected. Preference and index are configuration parameters.

Table 112 **Diameter Node Routing Table Field Descriptions (Continued)**

Field	Description
Server-Identifier	The next hop peer name

```
show aaa diameter node "node1.myrealm.com" peer "relay.telenet.be"
===============================================================================
Diameter Node : node1.myrealm.com
===============================================================================
Index                   : 1
Remote Host Identity    : relay.telenet.be
Remote Realm            : telenet.be
Remote IP address       : 13.0.1.5
Remote Port:  3868
Remote Origin-State-Id : 1522062441
Local Host Identity     : bng1.bng.nokia.com
Local Realm             : bng.nokia.com
Local IP address        : 13.0.1.4
Local port:  3456
Peer Admin State        : Up
Peer Active             : Yes
Peer Status             : I-Open
Active Applications     : nasreq gx
Last Disconnect Cause   : rebooting
Preference              : 4
Connection Timer (Tc)   : 10
Watchdog Timer   (Tw)   : 237
...
Last Mgmt Change        : 03/23/2018 11:42:43
```

Table 113 describes Diameter node peer fields.

Table 113 **Diameter Node Peer Field Descriptions**

Field	Description
Diameter Node	The Diameter node under which this peer is active
Index	The Peer index is a peer configuration parameter
Remote Host Identity	The peer identity
Remote Realm	The realm to which the peer belongs
Remote IP address	The IP address of the peer
Remote Port	The port on the peer used for Diameter communication
Remote Origin-State-Id	The Origin-State-Id of the peer
Local Host Identity	The local identity
Local Realm	The local realm name

Table 113 **Diameter Node Peer Field Descriptions (Continued)**

Field	Description
Local IP address	The local IP address
Local Port	The local port used for Diameter communication
Peer Admin State	The administrative state of the configured peer
Peer Active	Indicates whether the connection towards the peer is active
Peer Status	The operational status of the configured peer
Active Applications	The applications used over this peer
Last Disconnect Cause	The last disconnect cause
Preference	The configured preference of the peer
Connection Timer (Tc)	The timer used to initiate connections towards the peer. In case that the connection is not successfully established, a new attempt is periodically tried (every Tc interval)
Watchdog Timer (Tw)	The watchdog timer value that detects the failed connection towards the peer
Last Mgmt Change	The date and time of the last management change

```
show aaa diameter node "node1.myrealm.com" peer-table
===============================================================================
Diameter Peer Table (Node: node1.myrealm.com)
===============================================================================
Host Identity                       StatusT       Susp
-------------------------------------------------------------------------------
relay.telenet.be
                                    I-Open        No
relay.proximus.be
                                    I-Open        No
relay.proximus.be
                                    Closed        No
-------------------------------------------------------------------------------
No. of Active Peers: 2
```

Table 114 describes Diameter node peer table fields.

Table 114 **Diameter Node Peer Table Field Descriptions**

Field	Description
Node	The Diameter node name
Host Identity	The peer identity

Table 114 **Diameter Node Peer Table Field Descriptions (Continued)**

Field	Description
StatusT	The status of the peer
Susp	The suspension state of the peer (based on watchdogs)
Server-Identifier	The next hop peer name

```
show aaa diameter-node "gyrouter.workstation.be" peer "gyserver-
1.workstation.be" statistics
===============================================================================
Peer "gyserver-1.workstation.be"
===============================================================================
Message                                Sent                Received
-------------------------------------------------------------------------------
Capabilities-Exchange-Request          2                   0
Capabilities-Exchange-Answer           0                   2
Disconnect-Peer-Request                0                   0
Disconnect-Peer-Answer                 0                   0
Device-Watchdog-Request                188                 0
Device-Watchdog-Answer                 0                   188
Application message request            2367                0
Application message answer             0                   2366
Last cleared time: N/A
===============================================================================
```

diameter-peer-policy

Syntax **diameter-peer-policy** [*peer-policy-name*]
diameter-peer-policy *peer-policy-name* **peer** *peer-name* [**statistics**]
diameter-peer-policy *peer-policy-name* **client-side-peer-ip** *ip-address* **port** *tcp-port* [**statistics**]

Context show>aaa

Description This command displays diameter peer policy information and diameter peer information and statistics. For the diameter proxy, the incoming peer can be specified with an IP address and port number.

Parameters *peer-policy-name* — Specifies the diameter peer policy name, up to 32 characters.

peer-name — Specifies the diameter peer name, up to 32 characters.

ip-address — Specifies the IP address of the incoming peer (client side) of a diameter proxy.

tcp-port — Specifies the TCP port number of the incoming peer (client side) of a diameter proxy.

statistics — Displays the diameter peer statistics with details per diameter application, per message type, and per direction (rx/ tx).

Output The following is an example of AAA Diameter peer policy information.

Sample Output

```
# show aaa diameter-peer-policy "diameter-peer-policy-1"
===============================================================================
Diameter Peer Policy : diameter-peer-policy-1
===============================================================================
Last Mgmt Change      : 01/07/2016 20:05:12
Applications          : gy
Role                  : client
Description           : Diameter peer policy - gy
-------------------------------------------------------------------------------
Diameter Config Values
-------------------------------------------------------------------------------
Origin Host           : bng.domain.com
Origin Realm          : domain.com
Source IPv4 Address   : 192.0.2.51
Source IPv6 Address   : N/A
Router                : Base               Transaction Timer  : 30 (default)
Watchdog Timer        : 30 (default)       Connection Timer   : 30 (default)
Vendor Support        : 3GPP (default)
Python Policy         : py-policy-diam-1
-------------------------------------------------------------------------------
Peer Name            Oper  PSM State     Susp  Cooldown  Pref  Order  Pri/Sec
-------------------------------------------------------------------------------
peer-1               Yes   I-Open        No    -         30    1      Primary
peer-2               No    Closed        No    -         50    -      -
===============================================================================
# show aaa diameter-peer-policy "diameter-peer-policy-1" peer "peer-1"
===============================================================================
Diameter Peer Policy : diameter-peer-policy-1
===============================================================================
Peer IP address        : 172.16.1.1
Last Mgmt Change       : 01/07/2016 20:05:12
Peer Operational       : Yes              Watchdog Algorithm Active   : Yes
Peer State Machine     : I-Open           Watchdog Answer Pending     : No
Connection Timer  (Tc) : -                Connection Suspended        : No
Transaction Timer (Tt) : -                Cooldown Sequence Pending   : No
Watchdog Timer    (Tw) : 17               Cooldown Sequence Active    : No
Primary/Secondary Peer : Primary          Cooldown Sequence Progress  : -
Preference Order       : 1                Peer Removal Pending        : No
Pending Messages       : 0
===============================================================================
# show aaa diameter-peer-policy "diameter-peer-policy-1" peer "peer-1" statistics
===============================================================================
Diameter Peer Policy : diameter-peer-policy-1 (statistics)
===============================================================================
Diameter Peer             : peer-1
Time statistics cleared   : 01/07/2016 20:05:10
===============================================================================
                                    Request     Answer    Request     Answer
                                         Tx         Rx         Rx         Tx
===============================================================================
Aggregates
-------------------------------------------------------------------------------
Total                                     1          1          0          0
Failed                                    0          0          0          0
```

```
-------------------------------------------------------------------------------
Application ID 0 (Base)
-------------------------------------------------------------------------------
Capabilities-Exchange                    1          1          0          0
Disconnect-Peer                          0          0          0          0
Device-Watchdog                          0          0          0          0
-------------------------------------------------------------------------------
Application ID 1 (NASREQ)
-------------------------------------------------------------------------------
Authorization-Authentication             0          0          0          0
-------------------------------------------------------------------------------
Application ID 4 (Gy)
-------------------------------------------------------------------------------
Credit-Control (Initial)                 0          0          0          0
Credit-Control (Update)                  0          0          0          0
Credit-Control (Termination)             0          0          0          0
Re-Auth                                  0          0          0          0
Abort-Session                            0          0          0          0
-------------------------------------------------------------------------------
Application ID 16777238 (Gx)
-------------------------------------------------------------------------------
Credit-Control (Initial)                 0          0          0          0
Credit-Control (Update)                  0          0          0          0
Credit-Control (Termination)             0          0          0          0
Re-Auth                                  0          0          0          0
Abort-Session                            0          0          0          0
===============================================================================
```

Table 115 provides a description of the counters in the output of the **show aaa diameter-peer-policy** command. and Table 116 provides the failure reasons for the aggregate Failed counter (Request Rx and Answer Rx).

Table 115 **AAA Diameter Peer Policy Field Descriptions**

Counter		Counts
Aggregates Total	Request Tx	Aggregated counter — The sum of all successfully transmitted diameter request messages (all applications, all message types)
	Answer Rx	Aggregated counter — The sum of all successfully received diameter answer messages (all applications, all message types)
	Request Rx	Aggregated counter — The sum of all successfully received diameter request messages (all applications, all message types)
	Answers Tx	Aggregated counter — The sum of all successfully transmitted diameter answer messages (all applications, all message types)

Table 115 **AAA Diameter Peer Policy Field Descriptions (Continued)**

Counter		Counts (Continued)
Aggregates Failed	Request Tx	Aggregated counter — The sum of all request messages that could not be sent for internal reasons, such as TCP socket error (for example, congestion), failure reported to application (for example, no retries left) Dropped by Python — Timeouts (requests for which no answer was received within the configured timeout interval)
	Answer Rx	Aggregated counter — The sum of all Answer messages that failed to be successfully decoded, such as such as invalid header length or a violation against the Command Code Format (CCF) specification See Aggregate Failed Request Rx and Answer Rx counters Aggregate Failed Request Rx and Answer Rx counters for more information.
	Request Rx	Aggregated counter — The sum of all Request messages that failed to be successfully decoded, such as invalid header length or a violation against the Command Code Format (CCF) specification See Aggregate Failed Request Rx and Answer Rx counters Aggregate Failed Request Rx and Answer Rx counters for more information.
	Answers Tx	Aggregated counter — The sum of all answer messages that could not be sent for internal reasons such as TCP socket error (for example, congestion) Dropped by Python
Application ID *id* message type	Request Tx	Successfully transmitted diameter request messages for the corresponding application and message type
	Answer Rx	Successfully received diameter answer messages for the corresponding application and message type
	Request Rx	Successfully received diameter request messages for the corresponding application and message type
	Answers Tx	Successfully transmitted diameter answer messages for the corresponding application and message type

Table 116 **Aggregate Failed Request Rx and Answer Rx counters**

Role	Failed reasons
Proxy	Proxy (message length too short, too long, or version unsupported) Consistency check failed: header size, message length, illegal flag combination, AVP header too small, AVP present but not allowed (base RFC messages), and so on Not a base message Not a request Tx overload (message received to be send to client) Dropped by python Failed to create answer No transaction (answer to server, but request is not found - 5 min)
Diameter base / Diameter application	Message dropped due to result code Message too big / too small Dropped by python Consistency check failed end-2-end mismatch hop-to-hop command code (answer/request mismatch) Cool down sequence drop cc error request Application id in header error T bit in header

isa-radius-policy

Syntax **isa-radius-policy** *policy-name*
isa-radius-policy *policy-name* **associations**

Context show>aaa

Description This command displays ISA RADIUS policy information.

Parameters *policy-name* — Displays information about the specified ISA RADIUS policy.

associations — Displays the information associated with the ISA RADIUS server policy.

Output The following is an example of AAA ISA RADIUS policy information.

Sample Output

```
*B:asd-tr0610-dr421# show aaa isa-radius-policy "ZiggoAAA_DRP_ISAPlcy"
===============================================================================
Status for ISA RADIUS server policy "ZiggoAAA_DRP_ISAPlcy"
===============================================================================
Server 1, group 1, member 1
```

```
-------------------------------------------------------------------------------
Purposes Up                                  : accounting authentication
Source IP address                            : 172.18.128.33
Acct Tx Requests                             : 2469931
Acct Tx Retries                              : 320
Acct Tx Timeouts                             : 160
Acct Rx Replies                              : 2469471
Auth Tx Requests                             : 16417061
Auth Tx Retries                              : 7169
Auth Tx Timeouts                             : 2922
Auth Rx Replies                              : 16406973
CoA Rx Requests                              : 0
```

Table 117 describes subscriber ISA RADIUS policy output fields.

Table 117 **Subscriber ISA RADIUS Policy Field Descriptions**

Label	Description
Purposes Up	The RADIUS services that are up and running, and fully operational for this server
Source IP address	The IP address of the RADIUS server
Acct Tx Requests	The number of RADIUS transaction requests transmitted
Acct TX Retries	The number of RADIUS transaction request retries
Acct TX Timeouts	The number of RADIUS transaction requests that have timed out
Acct RX Replies	The number of RADIUS transaction responses received
Auth Tx Requests	The number of authentication requests transmitted
Auth Tx Retries	The number of authentication request retries
Auth Tx Timeouts	The number of RADIUS authentication requests that have timed out for the policy
CoA RX Requests	The number of Change-of-Authorization message responses received

radius-configuration

Syntax **radius-configuration**

Context show>aaa

Description This command displays RADIUS configuration information.

Output The following is an example of AAA RADIUS configuration information.

Sample Output

```
# show aaa radius-configuration
===============================================================================
RADIUS configuration
===============================================================================
CoA Port                    : 3799
===============================================================================
```

Table 117 provides a description of the counters in the output of the **show aaa isa-radius-policy** command.

Table 118 **AAA ISA RADIUS Policy Field Descriptions**

Label	Description
CoA Port	The RADIUS Change of Authorization (CoA) port

radius-server-policy

Syntax **radius-server-policy** *policy-name* [**acct-on-off**]
radius-server-policy *policy-name* **associations**
radius-server-policy *policy-name* **msg-buffer-stats**
radius-server-policy *policy-name* **statistics**
radius-server-policy [**acct-on-off**]

Context show>aaa

Description This command displays RADIUS server policy configuration information.

Parameters *policy-name* — Displays information for the specified RADIUS server policy.

association — Displays the information configured with the RADIUS server policy.

msg-buffer-stats — Displays statistics related to the RADIUS messages that are buffered for each specified RADIUS server policy.

statistics — Displays statistics for the specified RADIUS server policy.

act-on-off — Displays the admin state of the acct-on-off feature.

Output The following is an example of AAA RADIUS server policy information.

Sample Output

```
*B:asd-tr0610-dr421# show aaa radius-server-policy "ZiggoAAA_anycast" statistics
===============================================================================
RADIUS server policy "ZiggoAAA_anycast" statistics
===============================================================================
Tx transaction requests                         : 24818681
Rx transaction responses                        : 24817329
Transaction requests timed out                  : 1351
```

```
Transaction requests send failed                : 0
Packet retries                                  : 12410
Transaction requests send rejected              : 0
Authentication requests failed                  : 303530
Accounting requests failed                      : 0
Ratio of access-reject over auth responses      : 13%
Transaction success ratio                       : 99%
Transaction failure ratio                       : 1%
Statistics last reset at                        : 05/21/2015 01:11:39
```

Table 119 describes RADIUS server policy statistics output fields.

Table 119 **Subscriber RADIUS Server Policy Statistics Field Descriptions**

Label	Description
Tx transaction requests	The number of RADIUS transaction requests transmitted
Rx transaction responses	The number of RADIUS transaction responses received
Transaction requests timed out	The number of RADIUS transaction requests that have timed out
Transaction requests send failed	The number of RADIUS transaction requests that could not be transmitted
Packet retries	The number of times a RADIUS request packet was retransmitted to a server
Transaction requests send rejected	The number of RADIUS transaction requests that were not transmitted due to unacceptable configuration
Authentication requests failed	The number of authentication failures for this policy
Accounting requests failed	The number of accounting failures for this policy
Ratio of access-reject over auth responses	The ratio of access-rejects in the auth responses for this policy
Transaction success ratio	The transaction success ratio for this policy
Transaction failure ratio	The transaction failure ratio for this policy
Statistics last reset at	Indicated the date and time at which the statistics for this policy were last reset

wpp

Syntax **wpp portal-group** *name*
wpp [**portal-group** *name*] [**host** *ip-address*] **hosts**

Context show>aaa

Description This command displays information about the specified WPP portal group or the WPP hosts controlled by the WPP portal group.

Parameters *name* — Specifies the name of the WPP portal group up to 32 characters.

ip-address — Specifies the IP address of the WPP host.

hosts — Shows the status of triggered hosts.

Output The following is an example of WPP portal group or WPP host information.

Sample Output

```
show aaa wpp portal-group "test"
===============================================================================
wpp portal group "test"
===============================================================================
administrative state                  : in-service
controlled router                     : 500
number of enabled interfaces          : 1
triggered hosts                       : disabled
last management change                : 01/02/2017 19:02:11
===============================================================================
```

Table 120 describes WPP portal group output fields.

Table 120 **WPP Portal Group Field Descriptions**

Field	Description
administrative state	The administrative state of the WPP portal group
controlled router	The identifier of the virtual router instance being controlled by this portal
number of enabled interfaces	The number of interfaces associated with this portal that are administratively enabled
triggered hosts	Triggered hosts are enabled or disabled on this portal

ancp-policy

Syntax **ancp-policy** [*policy-name*]
ancp-policy *policy-name* **association**

Context show>subscr-mgmt

Description This command displays subscriber Access Node Control Protocol (ANCP) policy information.

Parameters *policy-name* — Displays information for the specified ANCP policy

association — Displays the information configured with the ANCP policy.

Output The following is an example of subscriber management ANCP policy information.

Sample Output

```
A:cses-E11>config>subscr-mgmt>ancp# show subscriber-mgmt ancp-policy "test"
===============================================================================
ANCP Policy "test"
===============================================================================
I. Rate Reduction      : 0 kbps
I. Rate Adjustment     : 100 percent
I. Rate Monitor        : 63360 kbps
I. Rate Monitor Alarm  : Yes
I. Rate Modify         : N/A

E. Rate Reduction      : 0 kbps
E. Rate Adjustment     : 100 percent
E. Rate Monitor        : 0 kbps
E. Rate Monitor Alarm  : no
E. Rate Modify         : N/A

Port Down : N/A

Last Mgmt Change: 02/13/2013 19:15:28
===============================================================================
*A:cses-E11>config>subscr-mgmt>ancp#
```

Table 121 describes ANCP policy output fields.

Table 121 **ANCP Policy Field Descriptions**

Field	Description
I. Rate Reduction	The ingress rate reduction to be applied for this subscriber in order to account for its video viewership
I. Rate Adjustment	The ingress rate adjustment
I. Rate Monitor	The ingress rate monitor
I. Rate Monitor Alarm	The ingress rate monitor alarm
I. Rate Modify	The ingress rate modification that is to be applied
E. Rate Reduction	The egress rate reduction to be applied for this subscriber in order to account for its video viewership
E. Rate Adjustment	The egress rate adjustment

Table 121 **ANCP Policy Field Descriptions (Continued)**

Field	Description
E. Rate Monitor	The egress rate monitor
E. Rate Monitor Alarm	The egress rate monitor alarm
E. Rate Modify	The egress-rate modification that is to be applied
Port Down	The status of the port down configuration
Last Mgmt Change	The date and time of the last management change

ancp-string

Syntax **ancp-string**
ancp-string *ancp-string*
ancp-string customer *customer-id* **site** *customer-site-name*
ancp-string sap *sap-id*

Context show>subscr-mgmt

Description This command displays subscriber Access Node Control Protocol (ANCP) string information.

Parameters *ancp-string* — Specifies an Access Node Control Protocol (ANCP) string up to 63 characters.

customer-id — Specifies an existing customer ID.

Values 1 to 2147483647

customer-site-name — Specifies an existing customer site name up to 32 characters.

sap-id — Displays ANCP string information for the specified SAP ID.

Values

<sap-id>	null	<port-id | bundle-id | bpgrp-id | lag-id | aps-id>
	dot1q	<port-id | bundle-id | bpgrp-id | lag-id | aps-id | pw-id>:qtag1
	qinq	<port-id | bundle-id | bpgrp-id | lag-id | pw-id>:qtag1.qtag2
	atm	<port-id | aps-id>[:vpi/vci | vpi | vpi1.vpi2 | cp.conn-prof-id]
		cp keyword
		conn-prof-id [1 to 8000]
	frame	<port-id | aps-id>:dlci
	cisco-hdlc	slot/mda/port.channel
	cem	slot/mda/port.channel
	ima-grp	<bundle-id>[:vpi/vci | vpi | vpi1.vpi2 | cp.conn-prof-id]

	cp	keyword
	conn-prof-id	[1 to 8000]
port-id	slot/mda/port[.channel]	
bundle-id	bundle-<type>-slot/mda.<bundle-num>	
	bundle	keyword
	type	ima | fr | ppp
	bundle-num	[1 to 336]
bpgrp-id	bpgrp-<type>-<bpgrp-num>	
	bpgrp	keyword
	type	ima | ppp
	bpgrp-num	[1 to 2000]
aps-id	aps-<group-id>[.channel]	
	aps	keyword
	group-id	[1 to 64]
ccag-id	ccag-<id>.<path-id>[cc-type]:<cc-id>	
	ccag	keyword
	id	[1 to 8]
	path-id	[a | b]
	cc-type	[.sap-net |.net-sap]
	cc-id	[0 to 4094]
eth-tunnel	eth-tunnel-<id>[:<eth-tun-sap-id>]	
	id	[1 to 1024]
	eth-tun-sap-id	[0 to 4094]
lag-id	lag-<id>	
	lag	keyword
	id	[1 to 800]
pw-id	pw-<id>	
	pw	keyword
	id	[1 to 10239]
qtag1	[0 to 4094]	
qtag2	[* | 0 to 4094]	
vpi	[0 to 4095] (NNI)	
	[0 to 255] (UNI)	
vci	[1 | 2 | 5 to 65535]	
dlci	[16 to 1022]	
tunnel-id	tunnel-<id>.<private | public>:<tag>	
	tunnel	keyword
id	[1 to 16]	
tag	[0 to 4094]	

Output The following is an example of subscriber management ANCP string information.

Sample Output

```
show subscriber-mgmt ancp-string "ANCP-0000003-0000001"
===============================================================================
ANCP-String "ANCP-0000003-0000001"
===============================================================================
Type       : SUB - "4AACAHCU74"
State      : Up                 Ancp Policy: N/A
I. Rate    : 129 kbps           E. Rate    : 130 kbps
Adj I. Rate: N/A                Adj E. Rate: N/A
Act I. Rate: N/A                Act E. Rate: N/A
Service Id : 50 (VPRN)
Group      : linux
Neighbor   : 10.0.0.2:34885
Persist Key: N/A
-------------------------------------------------------------------------------
Actual-Net-Data-Rate-Upstream              : 129 kbits/s
Actual-Net-Data-Rate-Downstream            : 130 kbits/s
Minimum-Net-Data-Rate-Upstream             : 131 kbits/s
Minimum-Net-Data-Rate-Downstream           : 132 kbits/s
Attainable-Net-Data-Rate-Upstream          : 133 kbits/s
Attainable-Net-Data-Rate-Downstream        : 134 kbits/s
Maximum-Net-Data-Rate-Upstream             : 135 kbits/s
Maximum-Net-Data-Rate-Downstream           : 136 kbits/s
Minimum-Net-Low-Power-Data-Rate-Upstream   : 137 kbits/s
Minimum-Net-Low-Power-Data-Rate-Downstream : 138 kbits/s
Maximum-Interleaving-Delay-Upstream        : 139 ms
Actual-Interleaving-Delay-Upstream         : 140 ms
Maximum-Interleaving-Delay-Downstream      : 141 ms
Actual-Interleaving-Delay-Downstream       : 142 ms
DSL-Line-State                             : 2 (IDLE)
Access-Loop-Encapsulation                  : 16909056 (0x01020300)
===============================================================================
```

Table 122 describes ANCP string output fields.

Table 122 **ANCP String Field Descriptions**

Field	Description
Type	The ANCP type
State	The administrative state of the policy
Ancp Policy	The ANCP policy name
I. Rate	The ingress rate
E. Rate	The egress rate
Adj I. Rate	The adjusted ingress rate
Service Id	The service ID

Table 122 **ANCP String Field Descriptions (Continued)**

Field	Description
Group	The group name
Neighbor	The neighbor ID
Persist Key	The key value that can be used to track this SLAAC host in the persistence file
Actual-Net-Data-Rate-Upstream	The transmitted actual data upstream rate
Actual-Net-Data-Rate-Downstream	The transmitted actual data downstream rate
Minimum-Net-Data-Rate-Upstream	The minimum data upstream rate
Mimimum-Net-Data-Rate-Downstream	The minimum data downstream rate
Attainable-Net-Date-Rate-Upstream	The attainable net data upstream rate
Attainable-Net-Data-Rate-Downstream	The attainable net data downstream rate
Maximum-Net-Data-Rate-Upstream	The maximum net data upstream rate
Maximum-Net-Data-Rate-Downstream	The maximum net data downstream rate
Maximum-Interleaving-Delay-Upstream	The maximum interleaving delay upstream rate
Actual-Interleaving-Delay-Upstream	The actual interleaving delay upstream rate
Maximum-Interleaving-Delay-Downstream	The actual interleaving delay downstream rate
Actual-Interleaving-Delay-Downstream	The actual interleaving delay downstream rate
DSL-Line-State	The DSL line state
Access-Loop-Encapsulation	The access loop encapsulation rate

authentication

Syntax **authentication** *policy-name* **association**
authentication [*policy-name*]
authentication [*policy-name*] **statistics**
authentication coa-statistics

Context show>subscr-mgmt

Description This command displays subscriber management RADIUS authentication policy information and statistics.

Parameters *policy-name* — Specifies the subscriber management RADIUS authentication policy name, up to 32 characters, for which information is requested.

association — Displays SAP, interface, local user database host, AA and L2TP associations of this policy.

coa-statistics — Displays the overall statistics for incoming RADIUS Change of Authorization (CoA) messages and Disconnect Messages. For dropped requests, a counter for different drop reasons is available.

statistics — Displays a list of policies with basic statistics (without specifying a policy name) or detailed statistics, including per-server statistics for the specified policy-name. These statistics apply only to the legacy RADIUS server configuration method where the servers are directly configured in the authentication policy.

Output The following is an example of subscriber management authentication information.

Sample Output

```
# show subscriber-mgmt authentication
===============================================================================
Authentication Policies
===============================================================================
Name                            Description
-------------------------------------------------------------------------------
auth-policy-1                   Radius auth policy - servers
auth-policy-2                   Radius auth policy - radius-server-policy
-------------------------------------------------------------------------------
Number of Authentication Policies : 2
===============================================================================


# show subscriber-mgmt authentication "auth-policy-2"
===============================================================================
Authentication Policy auth-policy-2
===============================================================================
Description          : Radius auth policy - radius-server-policy
Re-authentication    : Yes                 Username Format      : MAC Address
PPPoE Access Method  : PAP/CHAP            Username Mac-Format  : "aa:"
PPP-Username Oper    : None
PPP-Domain-Name      : N/A
Username Oper        : None
Domain-Name          : N/A
```

```
Acct-Stop-On-Fail    :
RADIUS Server Policy : "aaa-server-policy-1"
Fallback Action      : deny
Last Mgmt Change     : 06/24/2013 21:16:50
-------------------------------------------------------------------------------
Include Radius Attributes
-------------------------------------------------------------------------------
Remote Id            : Yes                 Circuit Id           : Yes
NAS Port Id          : Yes                 NAS Identifier       : Yes
PPPoE Service Name   : Yes                 DHCP Vendor Class Id : Yes
Access Loop Options  : Yes                 MAC Address          : Yes
NAS Port Prefix      : None                NAS Port Suffix      : None
NAS-Port-Type        : Yes (standard)      Acct Session Id      : Host
Calling Station Id   : Yes (sap-string)    Called Station Id    : Yes
Tunnel Server Attr   : Yes                 DHCP Options         : Yes
NAS Port             : Yes
NAS Port Bits Spec   : *3s*1m*4p*12o*12i
-------------------------------------------------------------------------------
Radius Servers
-------------------------------------------------------------------------------
Router               : management + Base   Source Address       : N/A
Access Algorithm     : Direct              Retry                : 3
Timeout (s)          : 5                   Hold down time (s)   : 30
-------------------------------------------------------------------------------
Index IP Address      Port  Pend-Req-Limit Out/Overload time (s) Oper State
-------------------------------------------------------------------------------
No Radius Servers configured.
-------------------------------------------------------------------------------
Accept Radius Attributes
-------------------------------------------------------------------------------
No Matching Entries
-------------------------------------------------------------------------------
Radius Script Policies
-------------------------------------------------------------------------------
Access-Request          : "N/A"
Access-Accept           : "N/A"
Change-of-Authorization : "N/A"
===============================================================================


# show subscriber-mgmt authentication "auth-policy-2" association
===============================================================================
Authentication Policy auth-policy-2
===============================================================================
-------------------------------------------------------------------------------
SAP Associations
-------------------------------------------------------------------------------
No associations found.
-------------------------------------------------------------------------------
Interface Associations
-------------------------------------------------------------------------------
Service-Id : 3000 (VPRN)
 - If Name : group-int-ws-1-1
-------------------------------------------------------------------------------
Local-User-Db PPPoE Host Associations
-------------------------------------------------------------------------------
Local-User-Db : ludb-1
 - Host : host-1
-------------------------------------------------------------------------------
```

```
Local-User-Db DHCP Host Associations
-------------------------------------------------------------------------------
Local-User-Db : ludb-1
 - Host : default
-------------------------------------------------------------------------------
Application Assurance Associations
-------------------------------------------------------------------------------
No associations found.
===============================================================================
No associated L2TP groups found.
No associated L2TP tunnels found.


# show subscriber-mgmt authentication statistics
===============================================================================
Authentication Policy Statistics
===============================================================================
Policy Name                       Subscr. Pkts   Subscr. Pkts   Subscr. Pkts
                                  Authenticated  Rejected       Rejected
                                                                Send Failed
-------------------------------------------------------------------------------
auth-policy-1                     0              0              0
auth-policy-2                     0              0              0
-------------------------------------------------------------------------------
Number of Authentication Policies : 2
===============================================================================


# show subscriber-mgmt authentication "auth-policy-1" statistics
===============================================================================
Authentication Policy Statistics
===============================================================================
-------------------------------------------------------------------------------
Policy name                              : auth-policy-1
subscriber packets authenticated         : 0
subscriber packets rejected              : 0
subscriber packets rejected send failed  : 0
-------------------------------------------------------------------------------
radius server     requests  requests  requests  requests   requests requests
idx IP-address    accepted  rejected  no reply  md5 failed pending  send failed
-------------------------------------------------------------------------------
1 172.16.1.1      0         0         0         0          0        0
-------------------------------------------------------------------------------
===============================================================================


# show subscriber-mgmt authentication coa-statistics
===============================================================================
Radius Notify Statistics    Change-Of-Authorization   Disconnect-Messages
===============================================================================
Requests Received             7                         10
Requests Accepted             5                         6
Requests Rejected             2                         4
Requests Dropped              0                         0
    No Auth Policy found      0                         0
    Invalid message           0                         0
    Out of resources          0                         0
    Authentication failure    0                         0
===============================================================================
```

Table 123 **Subscriber Authentication Field Descriptions**

Label	Description
Re-authentication	Whether re-authentication on IPOE renewal is enabled
Username Format	The user name format in RADIUS message
PPPoE Access Method	The PPPoE access method, none, padi, pap-chap
Username Mac-Format	The username is represented when contacting a RADIUS server.
PPP-Username Oper	The PPP user name operation
PPP-Domain-Name	The PPP domain name
Username Oper	The username
Acct-Stop-On-Fail	The failure condition a RADIUS accounting-stop message indicating the failure has to be sent
RADIUS Server Policy	The RADIUS server policy name
Fallback Action	The authentication action in case no RADIUS server is available
Lack Mgmt Change	The time of the most recent management-initiated change to this policy server
Remote Id	The value of the remote ID sub-option
NAS Port Id	NAS port ID generation enabled or disabled
PPPoE Service Name	The PPPoE service name
Access Loop Options	The generation of the DSL forum access loop characteristics RADIUS attributes enabled or disabled
NAS Port Prefix	The type of prefix added to the NAS-Port attribute if included in RADIUS authentication request messages
NAS-Port-Type	The kind of value to be put in the RADIUS NAS-Port-Type attribute if included in RADIUS authentication request messages
Calling Station Id	The string will be put in the RADIUS Calling-Station-Id attribute if included in RADIUS authentication request messages
Tunnel Server Attr	The tunnel-server-attrs RADIUS attribute enabled or disabled
NAS Port	The NAS-Port attribute included or excluded

Table 123 **Subscriber Authentication Field Descriptions (Continued)**

Label	Description (Continued)
NAS Port Bits Spec	The number that will be put in the RADIUS NAS-Port attribute if included in RADIUS authentication request messages
Circuit Id	Circuit ID generation enabled or disabled
NAS Identifier	NAS-Identifier RADIUS attribute enabled or disabled
DHCP Vendor Class Id	dhcp-vendor-class-id RADIUS attribute enabled or disabled
MAC Address	The generation of the client MAC address RADIUS attribute enabled or disabled
NAS Port Suffix	The type of suffix will be added to the NAS-Port-Id attribute if included in RADIUS accounting-request messages
Acct Session Id	The RADIUS accounting session id for this subscriber host
DHCP Options	DHCP options enabled or disabled
Router	The RADIUS router name
Access Algorithm	The algorithm used to access the set of RADIUS servers
Timeout (s)	The timeout for a response from the RADIUS server
Source Address	The source address of the RADIUS message
Retry	The number of retries for contacting the RADIUS server
Hold down time (s)	The hold down time
Index IP Address	The IP address of the index
Port	The port ID
Pend-Req-Limit	The limit of the number of pending RADIUS authentication requests for this RADIUS server
Out/Overload time (s)	The accumulated time that the value of the object RADIUS server has been equal to 'overloaded' since created
Oper State	The operational state
Access-Request	The number of Access-Request packets received by this RADIUS proxy server
Access-Accept	The number of Access-Accept packets received by this RADIUS proxy server
Change-of-Authorization	Change of authorization enabled or disabled

Table 123 **Subscriber Authentication Field Descriptions (Continued)**

Label	Description (Continued)
Policy Name	The authentication policy name
Subscr. Pkts Authenticated	The number of subscriber packets authenticated
Subscr. Pkts Rejected	The number of subscriber packets rejected
Subscr. Pkts Rejected Send Failed	The number of subscriber messages (such as DHCP and PPPoE) failed to be sent
subscriber packets authenticated	The number of subscriber packets authenticated
subscriber packets rejected	The number of subscriber packets rejected
subscriber packets send failed	The number of subscriber packets failed to be sent
Requests Received	The number of notify Change-of-Authorization requests received
Requests Accepted	The number of notify Change-of-Authorization requests accepted
Requests Rejected	The number of notify Change-of-Authorization requests rejected
Requests Dropped	The number of notify Change-of-Authorization requests dropped
No Auth Policy found	The number of notify Change-of-Authorization requests found
Invalid message	The number of notify Change-of-Authorization requests rejected because of decode errors
Out of resources	The number of notify Change-of-Authorization requests rejected due to lack of resources
Authentication Failure	The number of notify Change-of-Authorization requests which do not have NAS-Port-ID or Framed-IP-Address set or have mismatched subscriber-id

category-map

Syntax **category-map** [*category-map*]
category-map *category-map* **subscribers**

Context show>subscr-mgmt

Description This command displays category map information.

With no parameters specified, a summary view of all configured category maps is displayed.

With the optional *category-map* name parameter specified, the detailed configuration of the specified category map name is displayed.

With the *category-map* **subscribers** parameters specified, all subscribers that have the specified *category-map* associated are displayed.

Parameters *category-map* — Specifies the category map name for which information is displayed.

Output The following is an example of subscriber management category map information.

Sample Output

```
Node# show subscriber-mgmt category-map "cat-map-1"
===============================================================================
Category Map: cat-map-1
===============================================================================
Description        : gy - time and volume
Credit Type        : volume             Credit Exh. Thresh.: 80%
Activity Thresh.   : 0 kb/s
Gx session level
usage monitoring   : false
-------------------------------------------------------------------------------
Category: cat-1-time
-------------------------------------------------------------------------------
Description        : Rating group for Time quota
Credit Type Over.  : none               Rating Group       : 11
OutOfCredit Action : none
Default Credit     : none
Ingress Queues     : 1
Egress Queues      : 1
Ingress Policers   : none
Egress Policers    : none
ExhCrdSvcLvl Pir   : none
-------------------------------------------------------------------------------
Category: cat-1-volume
-------------------------------------------------------------------------------
Description        : Rating group for Volume quota
Credit Type Over.  : none               Rating Group       : 12
OutOfCredit Action : none
Default Credit     : none
Ingress Queues     : 1
Egress Queues      : 1
Ingress Policers   : none
Egress Policers    : none
ExhCrdSvcLvl Pir   : 100
===============================================================================
```

Table 124 describes category map output fields.

Table 124 **Category Map Field Descriptions**

Field	Description
Description	The user provided description of this policy
Credit Type	The volume or time based accounting is performed
Credit Exh. Thresh	The credit exhaust threshold taken into account to take action
Activity Thresh.	The threshold (in Kbps) that is applied to determine whether or not activity is occurring
Gx session level usage monitoring	Whether category map supports GX session level usage monitoring
Credit Type Over.	Credit type override
Rating Group	The rating group applicable for this category
OutOfCredit Action	The action to be taken if the credit es exhausted
Default Credit	The default value for the Time credit
Ingress Queues	The ingress queues defined in this category
Egress Queues	The egress queues defined in this category
Ingress Policers	The ingress policers defined in this category
Egress Policers	The egress queues defined in this category
ExhCrdSvcLvl Pir	The credit exhaust threshold taken into account to take action

diameter-application-policy

Syntax **diameter-application-policy** [*name*]

Context show>subscr-mgmt

Description This command displays Diameter application policy information.

Parameters *name* — Specifies the application policy, up to 32 characters, for which orphaned Gx sessions are displayed.

Output The following is an example of subscriber management Diameter application policy information

Sample Output

```
# show subscriber-mgmt diameter-application-policy
```

```
===============================================================================
DIAMETER application policies
===============================================================================
Name                            Description
-------------------------------------------------------------------------------
diameter-gx-policy-1            Diameter Gx policy
diameter-gy-policy-1            Diameter Gy policy
diameter-nasreq-policy-1        Diameter NASREQ policy
-------------------------------------------------------------------------------
No. of policies: 3
===============================================================================


# show subscriber-mgmt diameter-application-policy "diameter-nasreq-policy-1"
===============================================================================
DIAMETER application policy "diameter-nasreq-policy-1"
===============================================================================
Description                 : Diameter NASREQ policy
Session failover            : enabled
Failover handling           : continue
Peer policy                 : diameter-peer-policy-1
Application                 : nasreq
Tx timer (s)                : 10
Last management change      : 02/28/2015 14:53:49
-------------------------------------------------------------------------------
NASREQ
-------------------------------------------------------------------------------
Include AVP                 : nas-port-id
                              nas-port-type
NAS-Port-Id prefix type     : none
NAS-Port-Id suffix type     : user-string
NAS-Port-Id suffix          : @bng1
NAS-Port-Type type          : standard

User name format            : mac
User name operation         : no-operation
MAC address format          : aa:
Last management change      : 02/28/2015 14:53:49
===============================================================================
Interfaces using diameter-auth-policy "diameter-nasreq-policy-1"
-------------------------------------------------------------------------------
Interface-name                  Service-id Type
-------------------------------------------------------------------------------
group-int-1-1                   1000       IES
-------------------------------------------------------------------------------
No. of interfaces: 1
-------------------------------------------------------------------------------
VPLS SAP's with diameter-auth-policy "diameter-nasreq-policy-1"
-------------------------------------------------------------------------------
Service    SAP
-------------------------------------------------------------------------------
10         1/1/4:*.*
-------------------------------------------------------------------------------
No. of SAP's: 1
-------------------------------------------------------------------------------


*A:Dut-C# show subscriber-mgmt diameter-application-policy "diamapppol_gx"
===============================================================================
```

```
DIAMETER application policy "diamapppol_gx"
===============================================================================
Description                 : (Not Specified)
Session failover            : enabled
Failover handling           : retry-and-terminate
Peer policy                 : diampeerpol_gx
Application                 : gx
Tx timer (s)                : 10
Last management change      : 05/08/2015 05:55:59
-------------------------------------------------------------------------------
Gx
-------------------------------------------------------------------------------
Include AVP                 : an-gw-address
Calling-Station-Id type     : mac
NAS-Port bits spec          : 0
NAS-Port-Id prefix type     : user-string
NAS-Port-Id prefix          : Testing
NAS-Port-Id suffix type     : circuit-id
NAS-Port-Type value         : 0
User-Equipment-Info         : mac

Subscription-Id-Data origin : subscriber-id
Subscription-Id-Data type   : e164
MAC address format          : aa:
Report IP address event     : enabled
CCR-t replay interval       : 60
Last management change      : 05/08/2015 06:54:27
```

Table 125 describes diameter application policy output fields.

Table 125 **Diameter Application Policy Field Descriptions**

Field	Description
Description	The user provided description of this policy
Session failover	Failover is enabled or disabled
Failover Handling	Handling is retry-and-terminate or terminate
Peer policy	The diameter application peer policy name
Application	The policy's application, gx, gy, or nasreq
Tx timer (s)	The maximum wait time (seconds) while a client is pending
Last management change	The time of the most recent management-initiated change to this policy
Include AVP	The attribute-value-pairs for Gx messages
Calling-Station-Id type	The string that is put in the RADIUS Calling-Station-Id attribute if included in RADIUS authentication request messages

Table 125 **Diameter Application Policy Field Descriptions (Continued)**

Field	Description
NAS-Port bits spec	The number that is put in the RADIUS NAS-Port attribute if included in RADIUS authentication request messages
NAS-Port-Id prefix type	The type of prefix that is added to the NAS-Port-Id attribute if included in RADIUS accounting-request messages
NAS-Port-Id prefix	The user configurable string that is added as the prefix to the NAS-Port attribute if included in RADIUS accounting-request messages
NAS-Port-Id suffix type	The user configurable string that is added as the suffix to the NAS-Port attribute if included in RADIUS accounting-request messages
NAS-Port-Type value	The value that is put in the RADIUS NAS-Port-Type attribute if included in RADIUS accounting-request messages
User-Equipment-Info	The user equipment info, eui64, imeisv, mac, or modified-eui64

diameter-session

Syntax **diameter-session**

Context show>subscr-mgmt

Description This command enables the context to display diameter session information.

ccrt-replay

Syntax **ccrt-replay** [**session-id** *session-id*] [**diameter-application-policy** *name*]
ccrt-replay diameter-application-policy *name* **summary**
ccrt-replay summary

Context show>subscr-mgmt>diam-session

Description This command displays:

- information for all diameter sessions that are in CCR-t replay mode:
  - Diameter session ID
  - Diameter Application Policy
  - Replay Time left

- statistics per diameter application policy for diameter sessions in CCR-t replay mode or that were in CCR-t replay mode and were dropped or terminated:
  - Sessions in CCR-t replay mode: the number of Diameter Gx or Gy sessions from the specified diameter application policy that are currently in CCR-t replay mode
  - Replay sessions dropped (maximum life-time expired): the number of Diameter Gx or Gy sessions from the specified diameter application policy that were in CCR-t replay mode and got dropped because no CCA-t was received before **max-lifetime** expiration
  - Replay sessions dropped (to free resources): the number of Diameter Gx or Gy sessions from the specified diameter application policy that were in CCR-t replay mode and got dropped because the maximum number of active Diameter sessions for Gx or Gy was reached
  - Replay sessions dropped (cleared): the number of Diameter Gx or Gy sessions from the specified diameter application policy that were in CCR-t replay mode and got dropped because the operator issued a **clear subscriber-mgmt diameter-session ccrt-replay diameter-application-policy** *name* **sessions** command
  - Replay sessions terminated (CCA-T received): the number of Diameter Gx or Gy sessions from the specified diameter application policy that were in CCR-t replay mode and for which a CCA-t was received before **max-lifetime** expiration
  - Statistics last cleared time: timestamp when the statistics were last cleared or n/a if the statistics were never cleared

Parameters *name* — Specifies the application policy, up to 32 characters.

session-id — Specifies a diameter session ID.

summary — Displays summarized information about CCR-t replay.

Output The following is an example of subscriber management CCR-t replay information.

Sample Output

```
*A:Dut-C# show subscriber-mgmt diameter-session ccrt-replay
===============================================================================
Diameter Sessions in CCR-t Replay Mode
===============================================================================
Session-id                                                  Replay Time Left
    Diameter Application Policy
-------------------------------------------------------------------------------
router.workstation.be;1431089354;13
    diamapppol_gx                                           0d 21:21:46
bng.nokia.com;1537773225;4
    diam-gy-1                                               0d 00:59:16
-------------------------------------------------------------------------------
No. of Matching Entries: 2
===============================================================================
*A:Dut-C# show subscriber-mgmt diameter-session ccrt-replay session-id ro
===============================================================================
Diameter Sessions in CCR-t Replay Mode
===============================================================================
Session-id                                                  Replay Time Left
    Diameter Application Policy
-------------------------------------------------------------------------------
```

```
router.workstation.be;1431089354;13
    diamapppol_gx                                              0d 21:21:27
-------------------------------------------------------------------------------
No. of Matching Entries: 1
===============================================================================
*A:Dut-C# show subscriber-mgmt diameter-session ccrt-replay summary
===============================================================================
Diameter Sessions in CCR-t Replay Mode
===============================================================================
Total Count : 2
===============================================================================
*A:Dut-C# show subscriber-mgmt diameter-session ccrt-replay diameter-application-
policy "diam-gy-1"
===============================================================================
Diameter Sessions in CCR-T Replay Mode
===============================================================================
Session-id                                                 Replay Time Left
    Diameter Application Policy
-------------------------------------------------------------------------------
bng.nokia.com;1537773225;6
    diam-gy-1                                              0d 00:58:33
-------------------------------------------------------------------------------
No. of Matching Entries: 1
===============================================================================
===============================================================================
Diameter CCR-T Replay status for policy "diam-gy-1"
===============================================================================
Sessions in CCR-T replay mode                        : 1
Replay sessions dropped (maximum life-time expired)  : 0
Replay sessions dropped (to free resources)          : 0
Replay sessions dropped (cleared)                : 0
Replay sessions terminated (CCA-T received)          : 1
Statistics last cleared time                         : N/A
===============================================================================
```

Table 126 describes CCRT replay output fields.

Table 126 **CCR-T Replay Field Descriptions**

Field	Description
Session-id	The diameter session ID that is in CCR-t replay mode
Diameter Application Policy	The diameter application policy name to which the Gx or Gy session in CCR-t replay mode belongs
Replay Time Left	The remaining **max-lifetime** before the Diameter Gx or Gy session in CCR-t replay mode is dropped
No. of Matching Entries	The number of entries matching the search criteria
Sessions in CCR-T replay mode	The number of Diameter Gx or Gy sessions from the specified diameter application policy that are in CCR-t replay mode

Table 126 **CCR-T Replay Field Descriptions (Continued)**

Field	Description
Replay sessions dropped (maximum life-time expired)	The number of Diameter Gx or Gy sessions from the specified diameter application policy that were in CCR-t replay mode and got dropped because no CCA-t was received before max-lifetime expiration
Replay sessions dropped (to free resources)	The number of Diameter Gx or Gy sessions from the specified diameter application policy that were in CCR-t replay mode and got dropped because the maximum number of Diameter sessions was reached
Replay sessions dropped (cleared)	The number of Diameter Gx or Gy sessions from the specified diameter application policy that were in CCR-t replay mode and got dropped because the operator issued a **clear subscriber-mgmt diameter-session ccrt-replay diameter-application-policy** *name* **sessions** command
Replay sessions terminated (CCA-T received)	The number of Diameter Gx or Gy sessions from the specified diameter application policy that were in CCR-t replay mode and for which a CCA-t was received before max-lifetime expiration
Statistics last cleared time	The timestamp when the statistics were last cleared or n/a if the statistics were never cleared

mac

Syntax **mac all** | *ieee-address* | **unknown**

Context show>subscr-mgmt>errors

Description This command sorts all the subscriber errors by MAC address.

Parameters **all** — Shows all errors starting with the lowest MAC address.

ieee-address — Shows only the error of a specific MAC address.

unknown — Shows errors with unknown MAC addresses, such as for PPPoA subscribers.

Output The following is an example of subscriber management listed by MAC addresses.

Sample Output

```
*A:eng-BNG-2# show  subscriber-mgmt errors mac all
===============================================================================
Subscriber management errors
===============================================================================
-------------------------------------------------------------------------------
```

```
MAC       : 00:00:10:10:12:13
-------------------------------------------------------------------------------
Time      : 2016/06/23 17:35:46.0
Service   : 1000
SAP       : 1/1/20:841
Error     : No Offer from DHCP server after Discover from DHCP client
Extra     : DHCPv4
Time      : 2016/06/23 17:35:16.0
Service   : 1000
SAP       : 1/1/20:841
Error     : No Offer from DHCP server after Discover from DHCP client
Extra     : DHCPv4
===============================================================================
```

Table 127 describes subscriber management MAC error output fields.

Table 127 **Subscriber Management MAC Error Field Descriptions**

Field	Description
MAC	The MAC address associated with the error
Time	The time the error was reported
Service	The service ID
SAP	The SAP on which the error occurred
Error	The error that occurred on the MAC address
Extra	Extra information about the error that occurred

sap

Syntax **sap all** | *sap-id*

Context show>subscr-mgmt>errors

Description This command sorts all the subscriber errors by SAP.

Parameters **all** — Shows all errors starting with the lowest SAP index.

sap-id — Shows only the error of a specific SAP index.

Output The following is an example of subscriber management errors listed by SAP.

Sample Output

```
A:BNG-2# show  subscriber-mgmt errors sap all
===============================================================================
Subscriber management errors
===============================================================================
-------------------------------------------------------------------------------
```

```
SAP       : 1/1/20:841
-------------------------------------------------------------------------------
Time      : 2016/06/23 17:35:46.0
Service   : 1000
MAC       : 00:00:10:10:12:13
Error     : No Offer from DHCP server after Discover from DHCP client
Extra     : DHCPv4
Time      : 2016/06/23 17:35:16.0
Service   : 1000
MAC       : 00:00:10:10:12:13
Error     : No Offer from DHCP server after Discover from DHCP client
Extra     : DHCPv4
===============================================================================
```

Table 128 describes subscriber management SAP error output fields.

Table 128 **Subscriber Management SAP Error Field Descriptions**

Field	Description
SAP	The SAP ID associated with the error
Time	The time this error was reported to the subscriber management errors subsystem
Service	The service ID associated with the error
MAC	The MAC address associated with the error
Error	The error that occurred on the SAP
Extra	Extra information about the error that occurred

sdp

Syntax **sdp all** | *sdp-id:vc-id*

Context show>subscr-mgmt>errors

Description This command sorts all the subscriber errors by SDP.

Parameters **all** — Shows all errors beginning with the lowest SDP index.

sdp-id:vc-id — Shows only the error of a specific SDP index.

Output The following is an example of subscriber management errors listed by SDP.

Output Sample

```
<<NEED SCREEN OUTPUT EXAMPLE HERE>>
```

Table 129 describes subscriber management SDP error output fields.

Table 129 **Subscriber Management SDP Error Field Descriptions**

Field	Description
SDP	The SDP ID associated with the error
Time	The time this error was reported to the subscriber management errors subsystem
Service	The service ID associated with the error
MAC	The MAC address associated with the error
Error	The error that occurred on the SDP
Extra	Extra information about the error that occurred

unknown-origin

Syntax **unknown-origin**

Context show>subscr-mgmt>errors

Description This command lists errors of an unknown origin. Errors shown do have a corresponding SAP ID or a SDP ID.

Output The following is an example of subscriber management unknown origin errors information.

Output Sample

```
<<NEED SCREEN OUTPUT EXAMPLE HERE>>
```

explicit-subscriber-map

Syntax **explicit-subscriber-map**

Context show>subscr-mgmt

Description This command displays explicit subscriber mappings.

Output The following is an example of subscriber management explicit subscriber map information.

Sample Output

```
B:Dut-A>show>subscr-mgmt# explicit-subscriber-map
===============================================================================
Explicit Subscriber Map
===============================================================================
Key                              Sub profile
```

```
                                    SLA profile
-------------------------------------------------------------------------------
sub_ident_A_1                       sub_prof80
                                    sla_prof80
-------------------------------------------------------------------------------
Number of Explicit Subscriber Mappings : 1
===============================================================================
B:Dut-A>show>subscr-mgmt#
```

Table 130 describes subscriber management explicit subscriber map output fields.

Table 130 **Subscriber Management Explicit Subscriber Map Field Descriptions**

Field	Description
Key	The explicit subscriber mapping
Number of Explicit Subscriber Mappings	The number of configured explicit subscriber mappings

host-lockout-policy

Syntax **host-lockout-policy**
host-lockout-policy *policy-name* **association**
host-lockout-policy *policy-name*
host-lockout-policy *policy-name lockout-state* **sap** *sap-id* [**summary**]
host-lockout-policy *policy-name lockout-state* [**circuit-id** *circuit-id*] [**sap** *sap-id*] [**summary**]
host-lockout-policy *policy-name lockout-state* [**mac** *ieee-address*] [**sap** *sap-id*] [**summary**]
host-lockout-policy *policy-name lockout-state* [**remote-id** *remote-id*] [**sap** *sap-id*] [**summary**]

Context show>subscr-mgmt

Description This command displays host lockout policy information.

Parameters *policy-name* — Specifies a specific subscriber host lockout policy name up to 32 characters.

association — Specifies to display the association of the host lockout policy.

lockout-state — Specifies to display information for hosts in a specific lockout state.

Values all, active, graceful

sap *sap-id* — Specifies to display SAP ID information.

circuit-id *circuit-id* — Specifies to display circuit ID information. The *circuit-id* is 256 characters maximum.

mac *ieee-address* — Specifies to display MAC address information up to 30 characters.

remote-id *remote-id* — Specifies to display remote ID information up to 256 characters.

Output The following is an example of subscriber management host lockout policy information.

Sample Output

```
*A:cses-E11# show subscriber-mgmt host-lockout-policy
===============================================================================
Host Lockout Policies
===============================================================================
Lockout Policy                   Last Mgmt Change
  Lockout Time Min                 Lockout Time Max
Description
  Lockout Reset Time               Max Lockout Hosts
-------------------------------------------------------------------------------
test                             04/20/2012 19:51:02
  10                               3600
test
  60                               100
===============================================================================
*A:cses-E11#


*A:cses-E11# show subscriber-mgmt host-lockout-policy "test"
===============================================================================
Host Lockout Policy "test"
===============================================================================
Description                      test
Last Mgmt Change                 04/20/2012 19:51:02
Lockout time min                 10
Lockout time max                 3600
Lockout reset time               60
Max lockout hosts                100
Host key                         all
===============================================================================
*A:cses-E11#
```

Table 130 describes subscriber management host lockout policy output fields.

Table 131 **Subscriber Management Host Lockout Policy Field Descriptions**

Field	Description
Lockout Policy	The subscriber Host Lockout Policy name
Last Mgmt Change	The sysUpTime at the time of the most recent management-initiated change to this policy
Lockout Time Min	The minimum lockout-time for this host lockout policy
Lockout Time Max	The maximum lockout-time for this host lockout policy
Description	The user-provided description of this subscriber Host Lockout Policy

Table 131 **Subscriber Management Host Lockout Policy Field Descriptions (Continued)**

Field	Description
Max Lockout Hosts	The maximum number of lockouts for this host
Last Mgmt Change	The sysUpTime at the time of the most recent management-initiated change to this policy
Lockout time min	The minimum lockout-time for this Host Lockout Policy
Lockout time max	The maximum lockout-time for this Host Lockout Policy
Lockout reset time	The lockout reset time for this Host Lockout Policy
Max lockout hosts	The maximum number of locked out hosts for this Host Lockout Policy
Host key	The method for host matching

igmp-policy

Syntax **igmp-policy**
igmp-policy *policy-name* **association**
igmp-policy *policy-name*

Context show>subscr-mgmt

Description This command displays IGMP policy information and only applies to the 7750 SR.

Parameters *policy-name* — Specifies an existing IGMP policy.

association — Displays the information configured with the IGMP policy.

Output The following is an example of subscriber management IGMP policy information.

Sample Output

```
*B:Dut-C# show subscriber-mgmt igmp-policy
===============================================================================
IGMP Policies
===============================================================================
IGMP Policy
  Import Policy                    Admin Version
Description
  Num Subscribers                  Host Max Groups
  Fast Leave
-------------------------------------------------------------------------------
pol1
                                   3
  2                                0
  fast-leave
```

```
pol2
                                    3
  0                                 0
  fast-leave
===============================================================================
*B:Dut-C#


*B:Dut-C# show subscriber-mgmt igmp-policy "pol1"
===============================================================================
IGMP Policy pol1
===============================================================================
Import Policy                       :
Admin Version                       : 3
Num Subscribers                     : 2
Host Max Group                      : 0
Fast Leave                          : yes
===============================================================================
*B:Dut-C#


*B:Dut-C# show subscriber-mgmt igmp-policy "pol1" association
===============================================================================
IGMP Policy pol1 Associations
===============================================================================
sub_1
sub_2
-------------------------------------------------------------------------------
No. of subscriber(s): 2
===============================================================================
*B:Dut-C#
```

Table 132 describes subscriber management IGMP policy output fields.

Table 132 **Subscriber Management IGMP Policy Field Descriptions**

Field	Description
IGMP Policy	The IGMP policy name
Import Policy	The import policy name to filter packets
Admin Version	The version of IGMP which is running for this host
Description	The user-provided description of this IGMP policy
Num Subscribers	The number of subscribers
Host Max Groups	The maximum number of groups for which IGMP can have local receiver information based on received IGMP reports for this host
Fast Leave	Whether fast leave is allowed for this host
Pol1	Information about the first reported IGMP policy

Table 132 **Subscriber Management IGMP Policy Field Descriptions**

Field	Description
Pol2	Information about the second reported IGMP policy
Associations	The subscribers associated with the policy

ipoe-session-policy

Syntax **ipoe-session-policy** *ipoe-session-policy-name* [**association**]
ipoe-session-policy

Context show>subscr-mgmt

Description This command displays IPoE session policy information.

Parameters *ipoe-session-policy-name* — Specifies the IPoE session policy name up to 32 characters.

association — Displays the interface and captures SAPs that reference the IPoE session policy.

Output The following is an example of IPoE session policy information.

Sample Output

```
show subscriber-mgmt ipoe-session-policy "ipoe-policy-1"
===============================================================================
IPoE Session Policy "ipoe-policy-1"
===============================================================================
Description           : IPoE policy
Last Mgmt Change      : 02/28/2015 11:51:25
Session Key           : sap-mac
Session Timeout       : unlimited
===============================================================================


show subscriber-mgmt ipoe-session-policy "ipoe-policy-1" association
===============================================================================
IPoE Session Policy "ipoe-policy-1"
===============================================================================
-------------------------------------------------------------------------------
IPoE Interface Associations
-------------------------------------------------------------------------------
Service-Id : 1000 (IES)
 - group-int-1-1
Service-Id : 2000 (VPRN)
 - group-int-1-1
-------------------------------------------------------------------------------
Capture SAP Associations
-------------------------------------------------------------------------------
Service-Id : 10 (VPLS)
- 1/1/4:*.*
```

==========

Table 133 describes subscriber management IPoE session policy output fields.

Table 133 **Subscriber Management IPoE Session Policy Field Descriptions**

Field	Description
Description	The user-provided description of this IPoE policy
Last Mgmt Change	The sysUpTime at the time of the last modification
Session Key	The session key to be used for this IPoE session to group subscriber hosts
Session Timeout	The session timeout period to be used for this IPoE session
IPoE Interface Associations	The service ID and interface name of the associated interfaces
Capture SAP Associations	The service ID and SAP ID of the associated capture SAPS

local-user-db

Syntax **local-user-db** *local-user-db-name* **association** [**dhcp**] [**ppp**] [**12tp**] [**radius**] [**pppoe**] [**dhcp6**] [**capture-sap**] [**rtr-solicit**] [**wpp**] [**ipoe**]
local-user-db *local-user-db-name* **ipoe-all-hosts**
local-user-db *local-user-db-name* **ipoe-host** *ipoe-host-name*
local-user-db *local-user-db-name* **ipoe-unmatched-hosts**
local-user-db [*local-user-db-name*]
local-user-db *local-user-db-name* **pppoe-all-hosts**
local-user-db *local-user-db-name* **pppoe-host** *pppoe-host-name*
local-user-db *local-user-db-name* **pppoe-unmatched-hosts**

Context show>subscr-mgmt

Description This command displays local user database information.

Parameters *local-user-db-name* — Specifies the name of a local user database up, to 32 characters.

association — Displays entities associated with the specified local user database.

dhcp — Displays DHCP information.

ppp — Displays PPP information.

l2tp — Displays L2TP information.

radius — Displays RADIUS information.

pppoe — Displays PPPoE information.

dhcp6 — Displays DHCP6 information.

capture-sap — Displays capture SAP information.

rt-solicit — Displays RT solicit information.

wpp — Displays WPP information.

ipoe — Displays IPoE information.

ipoe-all-hosts — Displays IPoE all hosts information.

ipoe-host-name — Displays IPoE host name information.

ipoe-unmatched-hosts — Displays IPoE unmatched hosts information.

ppp-all-hosts — Displays information about PPP all hosts.

pppoe-hosts — Displays information about PPPoE hosts.

pppoe-unmatched-hosts — Displays information about PPPoE unmatched hosts.

Output The following is an example of local user database information.

Sample Output

```
*A:ALA-48>show>subscr-mgmt# local-user-db
===============================================================================
Local User Databases
===============================================================================
Name                             Admin Host  Description
                                 State Count
-------------------------------------------------------------------------------
database01                       Down  1
database02 Provider001/Class0002 Down  0     This is a long testdescription wi*
test                             Down  2
-------------------------------------------------------------------------------
Number of Local User Databases : 3     Number of Hosts : 3
===============================================================================
* indicates that the corresponding row element may have been truncated.


*A:ALA-48>show>subscr-mgmt# local-user-db database01
===============================================================================
Local User Database "database01"
===============================================================================
Admin State          : Down
Last Mgmt Change     : 11/08/2007 12:27:36
Host Count           : 1
DHCP Match Types     : circ-id
DHCP CircId Mask Pfx : test
DHCP CircId Mask Sfx : N/A
PPPoE Match Types    : N/A
PPPoE CircId Mask Pfx: N/A
PPPoE CircId Mask Sfx: N/A
===============================================================================
```

```
*A:ALA-48>show>subscr-mgmt#


*A:ALA-48>show>subscr-mgmt# local-user-db database01 dhcp-all-hosts
===============================================================================
Local User Database "database01" DHCP hosts
===============================================================================
Name                            Admin     Matched objects
                                State
-------------------------------------------------------------------------------
host001                         Down      -
-------------------------------------------------------------------------------
Number of DHCP Hosts : 1
===============================================================================

*A:ALA-48>show>subscr-mgmt# local-user-db "database01" dhcp-host host001
===============================================================================
DHCP Host "host001"
===============================================================================
Admin State          : Down
Last Mgmt Change     : 11/08/2007 12:13:42

Host Identification
 Circuit Id          : N/A
 Mac Address         : N/A
 Remote Id           : N/A
 Sap Id              : N/A
 Service Id          : N/A
 String              : N/A
 Option 60           : N/A
 System Id           : N/A

Matched Objects      : N/A

Address              : N/A

Identification Strings
 Subscriber Id       : N/A
 SLA Profile String  : N/A
 Sub Profile String  : N/A
 App Profile String  : N/A
 ANCP String         : N/A
 Inter Destination Id: N/A
===============================================================================

*A:ALA-48>show>subscr-mgmt# local-user-db "database01" dhcp-unmatched-hosts
===============================================================================
Local User Database "database01" DHCP unmatched hosts
===============================================================================
Name                            Reason      Duplicate Host
-------------------------------------------------------------------------------
host002                         No match    N/A
host003                         Duplicate   host001
host004                         No match    N/A
host005                         Duplicate   host001
-------------------------------------------------------------------------------
Number of DHCP Unmatched Hosts : 4
===============================================================================
*A:ALA-48>show>subscr-mgmt#
```

```
*A:ALA-48>show>subscr-mgmt# local-user-db "database01" association
===============================================================================
DHCP Servers where database01 is used
===============================================================================
Server-Name                     Router-Name
-------------------------------------------------------------------------------
dhcpS1                          vprn1000
-------------------------------------------------------------------------------
No. of Server(s): 1
===============================================================================
Interfaces where database01 is used for authentication
===============================================================================
Interface-Name                  Service-Id Type
-------------------------------------------------------------------------------
No. of Interface(s): 0
===============================================================================
*A:ALA-48>show>subscr-mgmt#


*A:ALA-48>show>subscr-mgmt# local-user-db "database01" association dhcp
===============================================================================
DHCP Servers where database01 is used
===============================================================================
Server-Name                     Router-Name
-------------------------------------------------------------------------------
dhcpS1                          vprn1000
-------------------------------------------------------------------------------
No. of Server(s): 1
===============================================================================
*A:ALA-48>show>subscr-mgmt#
===============================================================================

# show subscriber-mgmt local-user-db "ludb-1" association ipoe
===============================================================================
IPoE client interface associations for ludb-1
===============================================================================
Interface-Name                          Svc-Id     Type
-------------------------------------------------------------------------------
group-int-1-1                           1000       IES
group-int-1-1                           2000       VPRN
-------------------------------------------------------------------------------
No. of Interface(s): 2
===============================================================================
===============================================================================
Capture SAP associations for ludb-1
===============================================================================
SAP                             Svc-Id     Type   PPPoE PPP IPoE DHCP DHCP6 RS
-------------------------------------------------------------------------------
1/1/4:1202.*                    10         VPLS   y         y    y    y     y
1/1/4:*.*                       10         VPLS   y         y    y    y     y
-------------------------------------------------------------------------------
No. of SAP(s): 2
===============================================================================
```

Table 134 describer subscriber management local user database output fields.

Table 134 **Subscriber Management Local User Database Field Descriptions**

Field	Description
Name	The local user database name
Admin State	The administrative state of the local user database
Host Count	The number of hosts associated with the local user database
Description	The user-provided description of the local user database
Number of Local User Databases	The number of local user databases configured on the node
Number of Hosts	The number of hosts configured on the node
Last Mgmt Change	The sysUpTime at the time of the last modification
DHCP Match Types	The DHCP match type
DHCP CircId Mask Pfx	The DHCP circuit ID mask prefix
DHCP CircId Mask Sfx	The DHCP circuit ID mask suffix
PPPoE Match Types	The PPPoE match types; circuit-id, derived-id, dual-stack-remote-id, encap-tag-range, ip, mac, option60, remote-id, sap-id, service-id, string, system-id
PPPoE CircId Mask Pfx	The PPPoE circuit ID mask prefix
PPPoE CircId Mask Sfx	The PPPoE circuit ID mask suffix
Circuit Id	The circuit ID to match during host identification
Mac Address	The MAC address to matching during host identification
Remote Id	The remote ID from option 82 to match against
Sap Id	The SAP ID from the Nokia vendor specific sub-option in option 82 to match against
Service Id	The service ID to match during host identification
String	The string from the Nokia vendor specific sub-option in option 82 to match against
Option 60	The Vendor-Identifying Vendor Option to match against
System Id	The system ID from the Nokia vendor specific sub-option in option 82 to match against
Subscriber Id	The subscriber ID which is encoded in the identification strings

Table 134 **Subscriber Management Local User Database Field Descriptions (Continued)**

Field	Description
SLA Profile String	The SLA profile string which is encoded in the identification strings
Sub Profile String	The subscriber profile string which is encoded in the identification strings
App Profile String	The application profile string which is encoded in the identification strings
ANCP String	The ANCP string which is encoded in the identification strings
Inter Destination Id	The intermediate destination identifier which is encoded in the identification strings
Name	The unmatched host name
Reason	Indicates why this host is not matched
Duplicate Host	The name of the host that is already matched and that has the same values, configured for the objects that correspond with the match types, as this unmatched host
Server-Name	The service name of the unmatched host
Router-Name	The router name of the unmatched host
Interface-Name	The interface name of the unmatched host
Service-Id	The service ID that matches the host identification the Nokia vendor specific sub-option in option 82, or the service where the IPoE session is established
Interface-Name	The interface name associated with the local user database
Svc-Id Svc-Id	The service ID associated with the local user database
Type	The service type
SAP	The capture SAP associated with the local user database
PPPoE	PPPoE associated with the capture SAP; Y or N
PPP	PPPoE associated with the capture SAP; Y or N
IpoE	IPoE associated with the capture SAP; Y or N
DHCP	DHCP associated with the capture SAP; Y or N

Table 134 **Subscriber Management Local User Database Field Descriptions (Continued)**

Field	Description
DHCP6	DHCP6 associated with the capture SAP; Y or N
RS	RS associated with the capture SAP; Y or N
No. of SAP(s)	The total number of SAPs

msap-policy

Syntax **msap-policy** [*msap-policy-name* [**association**]]

Context show>subscr-mgmt

Description This command displays Managed SAP policy information.

Parameters *msap-policy-name* — Displays information about the specifies MSAP policy.

association — Displays information related to the specifies MSAP policy.

Output The following is an example of managed SAP policy information.

Sample Output

```
*A:ALA-48>show>subscr-mgmt# msap-policy
===============================================================================
Managed SAP Policies
===============================================================================
Name                             Num    Description
                                 MSAPs
-------------------------------------------------------------------------------
test                             0      (Not Specified)
test 1                           0      (Not Specified)
-------------------------------------------------------------------------------
Number of MSAP Policies : 2
Number of MSAPs         : 0
===============================================================================
*A:ALA-48>show>subscr-mgmt#
```

Table 135 describes subscriber management MSAP policy output fields.

Table 135 **Subscriber Management MSAP Policy Field Descriptions**

Field	Description
Name	The name of the MSAP policy
Num MSAPs	The number of MSAPs

Table 135 **Subscriber Management MSAP Policy Field Descriptions**

Field	Description
Description	The user-provided description of the MSAP policy
Number of MSAP Policies	The total number of MSAP policies
Number of MSAPs	The total number of MSAPs

pcc-rule

Syntax **pcc-rule**
pcc-rule monitoring-key *key* **detail**
pcc-rule rule-id *id* **detail**
pcc-rule rule-name *rule-name*
pcc-rule rule-name *rule-name* **detail**
pcc-rule summary
pcc-rule monitoring-key *key*

Context show>subscr-mgmt

Description This command displays a list of PCC rules and associated monitoring keys in the system.

Parameters *key* — Displays details about a specific monitoring-key.

id — Displays details about a specific PCC rule.

rule-name — Displays information about a specific PCC rule.

summary — Displays summarized information for active rules in the system.

key — Displays information about a specific monitoring-key.

detail — Displays detailed information.

Output The following is an example of subscriber management PCC rule information.

Sample Output

```
show subscriber-mgmt pcc-rule summary
===============================================================================
PCC Rules Summary
===============================================================================
Nbr Active PCC Rules    : 26 / 1023
Nbr Active Combinations
  IPv4 Filter           : 2 / 4095
  IPv6 Filter           : 0 / 4095
  Egress Qos            : 1 / 4095
  Ingress Qos           : 1 / 4095
===============================================================================
```

```
show subscriber-mgmt pcc-rule
===============================================================================
                        Id      Dir  ForwardAction         QosAction
-------------------------------------------------------------------------------
name      : RULE_egress_FC
monitorKey: -
                        29      egr  -                     fc
name      : RULE_egress_UM
monitorKey: um_RULE_egress_UM
                        34      egr  -                     monitor
name      : RULE_ingress_FC
monitorKey: -
                        37      ingr -                     fc
name      : RULE_ingress_UM
monitorKey: um_RULE_ingress_UM
                        50      ingr -                     monitor
name      : RULE_egress_DROP
monitorKey: -
                        28      egr  drop                  -
name      : RULE_ingress_RDR
monitorKey: -
                        49      ingr fwd nh4               -
name      : RULE_egress_UM_FC
monitorKey: um_RULE_egress_UM_FC
                        35      egr  -                     fc monitor
…
===============================================================================


show subscriber-mgmt pcc-rule rule-name "RULE_ingress_RATE_LIMIT_UM_FC_RDR" detail
===============================================================================
PCC Rules
===============================================================================
PCC rule name         : RULE_ingress_RATE_LIMIT_UM_FC_RDR
PCC rule id           : 47
Monitoring key        : um_RULE_ingress_RATE_LIMIT_UM_FC_RDR
Flow status           : Enabled
Nbr of Flows          : 1 (ingress)
HTTP-Redirect         : -
Next-Hop Redir. IPv4  : 10.10.10.10
Next-Hop Redir. IPv6  : -
QoS Ingr. CIR/PIR     : 1000 kbps / 2000 kbps
QoS Egr. CIR/PIR      : - / -
FC change             : h2
-------------------------------------------------------------------------------
Flows
-------------------------------------------------------------------------------
Src. IP  : any                                  Src. Port: -
Dst. IP  : 75.24.24.17/32                       Dst. Port: -
Protocol : 6                                    DSCP     : cp60
-------------------------------------------------------------------------------
===============================================================================


show service active-subscribers pcc-rule subscriber "1/1/3:1.1|00:00:00:00:00:01"
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
```

```
Subscriber 1/1/3:1.1|00:00:00:00:00:01 (subprof1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:1/1/3:1.1 - sla:sla1
-------------------------------------------------------------------------------
Ingr Qos Policy Override : 3:P2
Egr  Qos Policy Override : 2:P2
-------------------------------------------------------------------------------
IP Address
                MAC Address       PPPoE-SID Origin
-------------------------------------------------------
22.1.0.1
                00:00:00:00:00:01 N/A       DHCP
-------------------------------------------------------
Ingr Filter Override : 5:P4
Egr  Filter Override : 6:P5
=======================================================
Preference Rule Id   Rule Name
-------------------------------------------------------
0          28        RULE_egress_DROP
0          29        RULE_egress_FC
0          30        RULE_egress_RATE_LIMIT
0          31        RULE_egress_RATE_LIMIT_FC
0          32        RULE_egress_RATE_LIMIT_UM
0          33        RULE_egress_RATE_LIMIT_UM_FC
0          34        RULE_egress_UM
0          35        RULE_egress_UM_FC
0          36        RULE_ingress_DROP
0          37        RULE_ingress_FC
0          38        RULE_ingress_FC_HTTP
0          39        RULE_ingress_FC_RDR
0          40        RULE_ingress_HTTP
0          41        RULE_ingress_RATE_LIMIT
0          42        RULE_ingress_RATE_LIMIT_FC
0          43        RULE_ingress_RATE_LIMIT_FC_RDR
0          44        RULE_ingress_RATE_LIMIT_RDR
0          45        RULE_ingress_RATE_LIMIT_UM
0          46        RULE_ingress_RATE_LIMIT_UM_FC
0          47        RULE_ingress_RATE_LIMIT_UM_FC_RDR
0          48        RULE_ingress_RATE_LIMIT_UM_RDR
0          49        RULE_ingress_RDR
0          50        RULE_ingress_UM
0          51        RULE_ingress_UM_FC
0          52        RULE_ingress_UM_FC_RDR
0          53        RULE_ingress_UM_RDR
=======================================================
```

Table 136 describes subscriber management PCC rule output fields.

Table 136 **Subscriber Management PCC Rule Field Descriptions**

Field	Description
Nbr Active PCC Rules	The number of active PCC rules

Table 136 **Subscriber Management PCC Rule Field Descriptions**

Field	Description
Nbr Active Combinations	The number of active IP PCC rule combinations
IPv4 Filter	The number of IPv4 filter combinations
IPv6 Filter	The number of IPv6 filter combinations
Egress QoS	The number of active egress QoS PCC rule combinations
Ingress QoS	The number of active ingress QoS PCC rule combinations
Id	The PCC ID
Dir	The flow direction on which the PCC-rule acts
ForwardAction	The filter action of the PCC-rule
QosAction	The Qos action of the PCC-rule. The value account(3) is only relevant if the PCC-rule was received by the system as part of a RADIUS Subscriber Service attribute
name	The PCC rule name
monitorKey	The monitor key name that can be received in a PCC rule. An empty string (length 0) indicates that no monitor key is defined in the PCC rule
PCC rule name	The name of a PCC rule
PCC rule id	The PCC rule ID
Monitoring key	The monitoring key defined in this PCC-rule. This value is only taken into account if **monitor** is set in the QoS action.
Flow status	The PCC rule flow status; enabled or disabled
Nbr of Flows	The number of flow rules defined in this PCC rule
HTTP-Redirect	The URL to which all data flows that pertain to this PCC-rule must be redirected
Next-Hop Redir. IPv4	The IPv4 address to which all IPv4 data flows that pertain to this PCC-rule be redirected
Next-Hop Redir. IPv6	The IPv6 address to which all IPv6 data flows that pertain to this PCC-rule be redirected
QoS Ingr. CIP/PIR	The ingress QoS CIP and PIR values
QoS Egr. CIP/PIR	The egress QoS CIP and PIR values
FC change	The forwarding class (FC) that is assigned to all data flows that pertain to this PCC-rule

Table 136 **Subscriber Management PCC Rule Field Descriptions**

Field	Description
Src. IP	The address type of the source address
Src. Port	The port type of the source address
Dst. IP	The address type of the destination address
Protocol	The IPv4 protocol value, and for IPv6 the next header type used in the match criteria for this flow • 6 — TCP • 7 — UDP
DSCP	The DSCP to be matched for this flow
Subscriber	The subscriber SAP
SLA Profile Instance sap	The SLA profile instance SAP
Ingr QoS Policy Override	The ID of the overriding ingress Qos policy
Egr QoS Policy Override	The ID of the overriding egress Qos policy
IP Address	The subscriber IP address
MAC Address	The subscriber MAC address
PPPoE-SID	The PPPoE session ID of this subscriber host
Origin	The origin of this subscriber host
Preference	The preference with respect to other subscriber profiles
Rule ID	The PCC rule ID
Rule Name	The PCC rule name

radius-accounting-policy

Syntax **radius-accounting-policy** *name* **association**
radius-accounting-policy [*name*]
radius-accounting-policy *name* **statistics**

Context show>subscr-mgmt

Description This command displays RADIUS accounting policy information.

Parameters *name* — Specifies the RADIUS accounting policy name.

association — Displays parameters associated with this RADIUS accounting policy.

statistics — Displays statistics associated with this RADIUS accounting policy.

Output The following output displays subscriber management RADIUS accounting policy information.

Sample Output

```
*B:asd-tr0610-dr421# show subscriber-mgmt radius-accounting-policy "ZiggoAcct1813"
statistics
===============================================================================
Radius Accounting Policy ZiggoAcct1813 Statistics
===============================================================================
Tx Requests      : 36035966           Rx Responses  : 36035966
Request Timeouts : 0                  Send Retries  : 2713
Send Failed      : 0
-------------------------------------------------------------------------------
Radius Servers
-------------------------------------------------------------------------------
Index IP Address      Tx Reqs     Rx Resps    Req Timeouts Req Send Failed
-------------------------------------------------------------------------------
1     172.18.129.36   9012635     9011762     873          0
2     172.18.129.37   9004736     9003814     922          0
3     172.18.129.68   9010236     9009925     311          0
4     172.18.129.69   9011115     9010465     650          0
===============================================================================
```

Table 137 describes subscriber management RADIUS accounting policy output fields.

Table 137 **Subscriber Management RADIUS Accounting Policy Field Descriptions**

Label	Description
Tx Requests/TX Reqs	Displays the number of accounting requests transmitted for this policy.
Rx Responses/Rx Resps	Displays the number of accounting responses received for this policy.
Request Timeouts/ Req Timeouts	Displays the number of accounting requests which have timed out for this policy.
Send Retries	Displays the number of retries to a different server for a single accounting request for this policy.
Send Failed Req Send Failed	Displays how many accounting requests failed because the packet could not be sent out for this policy.

Table 137 **Subscriber Management RADIUS Accounting Policy Field Descriptions (Continued)**

Label	Description (Continued)
Radius Servers	Displays a table in which the statistics associated with this RADIUS accounting policy are broken down by individual RADIUS server. The table columns are: Index—displays the index number assigned to the RADIUS server. The index determines the sequence in which the servers are queried for authentication requests. Servers are queried in order from lowest to highest index. IP Address—the address of the RADIUS server. TX Reqs—see TX Requests in this table. Rx Resps—see RX Responses in this table. Req Timeouts—see Request Timeouts in this table. Req Send Failed—see Send Failed in this table.

sla-profile

Syntax **sla-profile** [*sla-profile-name* [**association**]]

Context show>subscr-mgmt

Description This command displays SLA profile information.

Parameters *sla-profile-name* — Specifies an existing SLA profile name.

association — Displays the information configured with the specified *sla-profile-name*.

Output The following is an example of subscriber management SLA profile information.

Sample Output

```
A:Dut-A# show subscriber-mgmt sla-profile
===============================================================================
SLA Profiles
===============================================================================
Name                              Description
-------------------------------------------------------------------------------
sla_default
sla_prof100_VOIP
sla_prof110_VOIP
sla_prof120_VOIP
sla_prof130_VOIP
sla_prof140_VOIP
sla_prof230_VOIP
sla_prof80
sla_prof80_VOIP
```

```
sla_prof81_VOIP
sla_prof90_VOIP
sla_profPC1
sla_profPC2
sla_profPC3
-------------------------------------------------------------------------------
Number of SLA Profiles : 14
===============================================================================
A:Dut-A#


A:Dut-A# show subscriber-mgmt sla-profile sla_prof100_VOIP
===============================================================================
SLA Profile sla_prof100_VOIP
===============================================================================
Host Limit           : 3 (Remove Oldest)
Ingress Qos-Policy   : 100                    Egress Qos-Policy : 100
Ingress Queuing Type : Service-queuing
Ingress Filter-Id    : N/A                    Egress Filter-Id  : N/A
Last Mgmt Change     : 07/10/2006 12:55:33
-------------------------------------------------------------------------------
Ingress Queue Overrides
-------------------------------------------------------------------------------
Queue Rate      CIR       HiPrio  CBS     MBS
-------------------------------------------------------------------------------
2     4000      -         -       -       -
3     2500      -         -       -       -

-------------------------------------------------------------------------------
Egress Queue Overrides
-------------------------------------------------------------------------------
Queue Rate      CIR       HiPrio  CBS     MBS
-------------------------------------------------------------------------------
2     4000      -         -       -       -
3     2500      -         -       -       -
===============================================================================
A:Dut-A#


A:Dut-A# show subscriber-mgmt sla-profile sla_prof100_VOIP association
===============================================================================
SLA Profile sla_prof100_VOIP
-------------------------------------------------------------------------------
SAP Default-Profile Associations
-------------------------------------------------------------------------------
No associations found.
-------------------------------------------------------------------------------
SAP Static Host Associations
-------------------------------------------------------------------------------
No associations found.
-------------------------------------------------------------------------------
SAP Non-Sub-Traffic-Profile Associations
-------------------------------------------------------------------------------
No associations found.
-------------------------------------------------------------------------------
Sub-Ident-Policy Profile Map Associations
-------------------------------------------------------------------------------
Policy-name : sub_ident_all
 - Key : sla_prof100_VOIP
```

```
-------------------------------------------------------------------------------
Sub-Profile Map Associations
-------------------------------------------------------------------------------
No associations found.
-------------------------------------------------------------------------------
Explicit Subscriber Map Associations
-------------------------------------------------------------------------------
No associations found.
===============================================================================
A:Dut-A#
```

Table 138 describes subscriber management SLA profile output fields.

Table 138 **Subscriber Management SLA Profile Field Descriptions**

Field	Description
Name	The SLA profile name
Description	The user provided description of this policy
Number of SLA Profiles	The total number of SLA profiles configured on this node
Host Limit	The maximum number of hosts using a given SLA (Service Level Agreement) profile
Ingress Qos-Policy	The ingress QoS policy ID
Egress Qos-Policy	The egress QoS policy ID
Ingress Queuing Type	The queuing type
Ingress Filter-Id	The ingress IP filter
Egress Filter-Id	The egress IP filter
Last Mgmt Change	The time and date of the most recent management-initiated change
Queue	The ingress queue ID
Rate	The SLA profile queue rate
CIR	The CIR rate
HiPrio	The number of high priority packets as determined by the SLA profile ingress QoS policy
CBS	The amount of reserved buffer space (in bytes) for the policer
MBS	The amount of buffer space (in bytes) allowed for the policer

vport

Syntax **vport** *port-id* **vport** *name* [**scheduler** *scheduler-name*]

Context show>qos>scheduler-stats

Description This command displays the Vport scheduler stats.

Parameters *port-id* — Displays information for the specified port.

name — Displays information for the specified Vport.

scheduler-name — Displays information for the specified scheduler-name.

Output The following is an example of Vport scheduler information.

Sample Output

```
*A:BNG# show qos scheduler-stats port 1/1/1 vport "dslam1" scheduler "dslam-sched"
===============================================================================
Scheduler Stats
===============================================================================
Scheduler                          Forwarded Packets      Forwarded Octets
-------------------------------------------------------------------------------
Egress Schedulers

dslam-sched                        0                      0
===============================================================================
*A:BNG#
```

Table 139 describes Vport scheduler statistics output fields.

Table 139 **Vport Scheduler Statistics Field Descriptions**

Field	Description
Scheduler	The scheduler name
Forwarded Packets	The number of forwarded packets, as determined by the subscriber ingress scheduler policy
Forwarded Octets	The number of forwarded octets, as determined by the subscriber ingress scheduler policy

statistics

Syntax **statistics iom** (*slot* | **all**) [**host** | **session** | **subscriber** | **summary**] [**non-zero-value-only**]
statistics mda (*mda* | **all**) [**host** | **session** | **subscriber** | **summary**] [**non-zero-value-only**]
statistics port (*port-id* | **all**) [**host** | **session** | **subscriber** | **summary**] [**non-zero-value-only**]
statistics pw-port (*pw-port* | **all**) [**host** | **session** | **subscriber** | **summary**] [**non-zero-value-only**]

statistics sla-profiles
statistics sub-profiles
statistics system [**host** | **session** | **subscriber** | **summary**] [**non-zero-value-only**]

Context show>subscr-mgmt

Description This command displays enhanced subscriber management statistics per port, PW port, MDA, IOM, or system.

For each statistic, there is current value and peak value. Peak value is the highest value since the last reset as a result of a system boot or the **clear subscriber-mgmt peakvalue-stats** command.

Note: The peak values can be reset using the **clear subscriber-mgmt peakvalue-stats** command.

Parameters **iom** *slot* — Displays specified IOM slot information.

mda *mda* — Displays specified slot/mda information.

port *port-id* — Specifies to display information for both the physical port ID and LAG.

pw-port *pw-port* — Specifies to display information for a pseudowire port ID.

Values 1 to 10239

all — displays statistics of all IOM or MDA or port or pseudowire port in the system

host — Displays v4/v6 host statistics only.

session — Displays PPPoX/LAC/LNS session statistics only.

sla-profiles — Displays the current and historical statistics for SLA profiles.

sub-profiles — Displays the current and historical statistics for subscriber profiles.

subscriber — Displays subscriber statistics only.

summary — Displays summary statistics only.

non-zero-value-only — Displays only non-zero value counters.

The following tables describe the counters available in the **show subscriber management statistics** command output.

The following terminology is used to indicate applicability of the stats:

- ESM — Enhanced Subscriber Management. Subscriber traffic forwarded via subscriber queues. Enabled with SAP **sub-sla-mgmt** in no shutdown state
- BSM — Basic Subscriber Management. Subscriber traffic forwarded via SAP queues. SAP sub-sla-mgmt must be in shutdown state. For DHCP, dhcp lease-populate or dhcp6-relay lease-populate must be enabled to count the leases. For IPv4, if anti-spoof is enabled on the SAP, a subscriber host is instantiated.
- Routed CO — IES or VPRN service with subscriber interface and group interface constructs

- Bridged CO — VPLS service with DHCPv4 lease management enabled (lease-populate)
- regular interface — IES or VPRN interface (none subscriber-interface or group-interface)
- Host (also subscriber host) — A resource in the system that is used for traffic forwarding and security related actions. The creation of a subscriber host entry is linked to anti-spoof being enabled on a SAP. For ESM, anti-spoof is mandatory and hence every connected {IP/MAC} consumes by default a subscriber host entry. A DHCP6 IA-PD can also be modeled as a managed route. In this case, no subscriber host is instantiated. For BSM, anti-spoof is optional on regular interfaces. An IPv4 static-host and DHCPv4 lease do not result in a subscriber host instantiation when anti-spoof is disabled on the SAP.

Table 140 **Host And Protocol Statistics**

Section	Counter	Counts	Applies to
IPv4	1. PPP Hosts - IPCP	IPv4 local terminated PPP hosts (PTA, LNS)	ESM, Routed CO
	2. IPOE Hosts - DHCP	DHCPv4 hosts (lease states)	ESM, Routed CO, Bridged CO
	3. IPOE Hosts - ARP	ARP hosts	ESM, Routed CO, Bridged CO
	4. IPOE Hosts – Static	IPv4 static hosts	ESM, Routed CO, Bridged CO
	5. IPOE Hosts BSM - DHCP	DHCPv4 hosts (lease states: anti-spoof and lease-populate enabled)	BSM, Routed CO, Bridged CO, regular interface
	6. IPOE Hosts BSM – Static	IPv4 static hosts (with anti-spoof enabled)	BSM, Routed CO, Bridged CO, regular interface
	7. IPOE BSM - DHCP	DHCPv4 lease states (with lease-populate enabled, no anti-spoof)	BSM, Routed CO, Bridged CO, regular interface
	8. IPOE BSM – Static	IPv4 static hosts (no anti-spoof)	BSM, Routed CO, Bridged CO, regular interface

Table 140 **Host And Protocol Statistics (Continued)**

Section	Counter	Counts	Applies to
IPv6	PPP Hosts – SLAAC	Local terminated IPv6 wan-host – SLAAC (PTA, LNS)	ESM, Routed CO
	PPP Hosts - DHCP6 (PD)	Local terminated IPv6 pd-host (PTA, LNS) – DHCP6 IA-PD leases over PPP (excluding PD as managed route)	ESM, Routed CO
	PPP Hosts - DHCP6 (NA)	Local terminated IPv6 wan-host (PTA, LNS) – DHCP6 IA-NA leases over PPP	ESM, Routed CO
	PPP Mngd Rt - DHCP6 (PD)	IPv6 (PTA, LNS) – DHCP6 IA-PD leases over PPP (PD as managed route only)	ESM, Routed CO
	IPOE Hosts – SLAAC	IPv6 wan-host – SLAAC	ESM, Routed CO
	IPOE Hosts - DHCP6 (PD)	IPv6 pd-host – DHCP6 IA-PD leases (excluding PD as managed route)	ESM, Routed CO
	IPOE Hosts - DHCP6 (NA)	IPv6 wan-host – DHCP6 IA-NA leases	ESM, Routed CO
	IPOE Mngd Rt - DHCP6 (PD)	IPv6 – DHCP6 IA-PD leases (PD as managed route only)	ESM, Routed CO
	IPOE Hosts – Static (PD)	IPv6 static hosts with prefix-length shorter than /128	ESM, Routed CO
	IPOE Hosts – Static (WAN)	IPv6 static hosts with prefix-length equal to /128	ESM, Routed CO
	IPOE BSM - DHCP6 (PD)	IPv6 – DHCP6 IA-PD leases (lease-populate)	BSM, regular interface
	IPOE BSM - DHCP6 (NA)	IPv6 – DHCP6 IA-NA leases (lease-populate)	BSM, regular interface

Table 140 **Host And Protocol Statistics (Continued)**

Section	Counter	Counts	Applies to
Total	PPP Hosts	Local terminated PPP hosts (PTA, LNS) Sum of counters 1, 9, 10 and 11	ESM
	IPOE Hosts	Total IPv4 and IPv6 IPOE hosts. Sum of counters 2, 3, 4, 5, 6, 13, 14, 15, 17 and 18	ESM
	IPv4 Hosts	Total IPv4 hosts. PPP (PTA, LNS) and IPOE. Sum of counters 1, 2, 3, 4, 5 and 6	ESM
	IPv6 Hosts	Total IPv6 hosts. PPP (PTA, LNS) and IPOE. Sum of counters 9, 10, 11, 13, 14, 15, 17 and 18	ESM
	IPv6 PD Mngd Routes	Total DHCP6 IA-PD leases modeled as a managed route. PPP (PTA, LNS) and IPOE. Sum of counters 12 and 16	ESM
	L2TP LAC Hosts	L2TP LAC hosts – single host per single or dual stack PPP session. Counter also increases for outgoing LTS sessions.	ESM, Routed CO
	Internal Hosts	Subscriber hosts for internal use. For example: LNS redirect hosts (for LTS, an LNS redirect host is also instantiated).	ESM

Table 140 **Host And Protocol Statistics (Continued)**

Section	Counter	Counts	Applies to
Total (Cont)	Non-Sub-Traffic L2-Hosts	Host on a single subscriber SAP in a VPLS service that enables non-IP traffic to be forwarded using the specified SLA profile instance queues. Host on a single subscriber SAP attached to an IES/ VPRN group-interface that enables traffic normally forwarded via the SAP queues to flow via the specified SLA profile instance queues. configure service vpls <service-id> sap <sap-id> sub-sla-mgmt single-sub-parameters non-sub-traffic sub-profile <sub-profile-name> sla-profile <sla-profile-name> [subscriber <sub-ident-string>] [app-profile <app-profile-name>]	ESM, Routed CO, Bridged CO
	DHCP leases	Total number of DHCPv4 lease states. Sum of counters 2, 5 and 7	ESM, BSM
	DHCPv6 leases	Total number of DHCPv6 lease states. Sum of counters 10, 11, 12, 14, 15, 16, 19 and 20	ESM, BSM
	Subscriber Hosts	Counter displayed in the output of “show subscriber-mgmt statistics iom | mda | port | pw-port” This counter matches the number of hosts accounted for in the per line card limit Sum of counters 1, 2, 3, 4, 5, 6, 9, 10, 11, 13, 14, 15, 17, 18 and 26	ESM

Table 140 **Host And Protocol Statistics (Continued)**

Section	Counter	Counts	Applies to
Total (Cont)	System Hosts Scale	Counter displayed in the output of “show subscriber-mgmt statistics system” This counter matches the number of hosts accounted for in the system wide limit Sum of counters 1, 2, 3, 4, 5, 6, 9, 10, 11, 13, 14, 15, 17, 18, 26 and 27	ESM

Table 141 **PPP Session Statistics**

Section	Counter	Counts	Applies to
Local	PPP Sessions - PPPoE	Local terminated PPPoE sessions (PTA)	ESM, Routed CO
	PPP Sessions - PPPoEoA	Local terminated PPPoEoA sessions (PTA)	ESM, Routed CO
	PPP Sessions - PPPoA	Local terminated PPPoA sessions (PTA)	ESM, Routed CO
	PPP Sessions - L2TP (LNS)	Local terminated PPP sessions (L2TP LNS)	ESM, Routed CO
LAC	PPP Sessions - PPPoE	Tunneled PPPoE session (L2TP LAC)	ESM, Routed CO
	PPP Sessions - PPPoEoA	Tunneled PPPoEoA session (L2TP LAC)	ESM, Routed CO
	PPP Sessions - PPPoA	Tunneled PPPoA session (L2TP LAC)	ESM, Routed CO
	PPP Sessions - L2TP (LTS)	Tunneled PPP session (L2TP LTS)	ESM, Routed CO

Table 141 **PPP Session Statistics (Continued)**

Section	Counter	Counts	Applies to
Total	PPP Sessions - established	PPP sessions that are established (at least one active host attached) – PTA/LAC/ LTS/LNS	ESM, Routed CO
	PPP Sessions - in setup	PPP sessions in setup (session created, host setup in progress) – PTA/LAC/ LTS/LNS	ESM, Routed CO
	PPP Sessions - local	Local terminated PPPoX sessions (PTA, L2TP LNS) Sum of counters 33, 34, 35 and 36	ESM, Routed CO
	PPP Sessions - LAC	Tunneled PPPoX session (L2TP LAC, L2TP LTS) Sum of counters 37, 38, 39 and 40	ESM, Routed CO
L2TP	L2TP Tunnels - originator	Number of L2TP Tunnels originated on this node. (LAC/ LTS)	ESM, Routed CO
	.L2TP Tunnels - receiver	Number of L2TP Tunnels terminated on this node. (LNS/ LTS)	ESM, Routed CO
	Total L2TP Tunnels	Number of L2TP Tunnels originated or terminated on this node Sum of counters 45 and 46	ESM, Routed CO

Table 142 **IPoE Session Statistics**

Section	Counter	Counts	Applies to
Total	IPOE Sessions - established	IPoE sessions that are established (at least one active host attached).	ESM, Routed CO
	IPOE Sessions- in setup	IPoE sessions in setup (session created, host setup in progress).	ESM, Routed CO

Table 143 **Subscriber Statistics**

Section	Counter	Counts	Applies to
Total	Subscribers	Total number of active subscribers.	ESM, Routed CO, Bridged CO

Table 144 **SubMgmt Statistics Summary**

Section	Counter	Counts
Hosts	IPv4 IPv6	Total IPv4 hosts (counter 23 in tables above) Total IPv6 hosts (counter 24 in tables above)
Sessions	PPP IPOE	Total PPP sessions - established (counter 41 in tables above) Total IPOE sessions – established (counter 48 in tables above)
Subscribers		Total number of active subscribers (counter 50 in tables above)

Output The following is an example of subscriber management system statistics information.

Sample Output

```
A:PE-1# show subscriber-mgmt statistics system
===============================================================================
Subscriber Management Statistics for System
===============================================================================
       Type                                Current      Peak      Peak Timestamp
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Host & Protocol Statistics
-------------------------------------------------------------------------------
IPv4   PPP Hosts      - IPCP                   1         1 02/28/2015 16:25:43
       IPOE Hosts     - DHCP                   0         2 02/28/2015 12:38:58
       IPOE Hosts     - ARP                    1         1 02/28/2015 13:46:10
       IPOE Hosts     - Static                 0         0
       IPOE Hosts BSM - DHCP                   0         0
```

```
       IPOE Hosts BSM - Static                0         0
       IPOE BSM       - DHCP                  0         0
       IPOE BSM       - Static                0         0
-------------------------------------------------------------------------------
IPv6   PPP Hosts      - SLAAC                 0         0
       PPP Hosts      - DHCP6 (PD)            0         0
       PPP Hosts      - DHCP6 (NA)            0         0
       PPP Mngd Rt    - DHCP6 (PD)            0         0
       IPOE Hosts     - SLAAC                 0         0
       IPOE Hosts     - DHCP6 (PD)            0         0
       IPOE Hosts     - DHCP6 (NA)            0         0
       IPOE Mngd Rt   - DHCP6 (PD)            0         0
       IPOE Hosts     - Static (PD)           0         0
       IPOE Hosts     - Static (WAN)          0         0
       IPOE BSM       - DHCP6 (PD)            0         0
       IPOE BSM       - DHCP6 (NA)            0         0
-------------------------------------------------------------------------------
Total  PPP Hosts                              1         1 02/28/2015 16:25:43
       IPOE Hosts                             1         2 02/28/2015 12:38:58
       IPv4 Hosts                             2         2 02/28/2015 16:25:43
       IPv6 Hosts                             0         0
       IPv6 PD Mngd Routes                    0         0
       L2TP LAC Hosts                         0         0
       Internal Hosts                         0         0
       Non-Sub-Traffic L2-Hosts               0         0
       DHCP Leases                            0         2 02/28/2015 12:38:58
       DHCPv6 Leases                          0         0
       System Hosts Scale                     2         2 02/28/2015 16:25:43
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
PPP Session Statistics
-------------------------------------------------------------------------------
Local  PPP Sessions  - PPPoE                  1         1 02/28/2015 16:25:43
       PPP Sessions  - PPPoEoA                0         0
       PPP Sessions  - PPPoA                  0         0
       PPP Sessions  - L2TP (LNS)             0         0
-------------------------------------------------------------------------------
LAC    PPP Sessions  - PPPoE                  0         0
       PPP Sessions  - PPPoEoA                0         0
       PPP Sessions  - PPPoA                  0         0
       PPP Sessions  - L2TP (LTS)             0         0
-------------------------------------------------------------------------------
Total  PPP Sessions  - established            1         1 02/28/2015 16:25:43
       PPP Sessions  - in setup               0         1 02/28/2015 16:25:43
       PPP Sessions  - local                  1         1 02/28/2015 16:25:43
       PPP Sessions  - LAC                    0         0
-------------------------------------------------------------------------------
L2TP   L2TP Tunnels  - originator             0         0
       L2TP Tunnels  - receiver               0         0
       Total L2TP Tunnels                     0         0
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
IPOE Session Statistics
-------------------------------------------------------------------------------
Total  IPOE Sessions - established            0         0
       IPOE Sessions - in setup               0         0
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Subscriber Statistics
```

```
-------------------------------------------------------------------------------
Total  Subscribers                              2         2 02/28/2015 16:25:43
-------------------------------------------------------------------------------
===============================================================================
Peak values last reset at : n/a
```

Sample Output (summary view)

```
A:PE-1# show subscriber-mgmt statistics port 1/1/4 summary
===============================================================================
SubMgmt Statistics
===============================================================================
                 |       Hosts         |      Sessions      |    Subscribers
Port Id          |   IPv4      IPv6    |    PPP      IPOE   |
-------------------------------------------------------------------------------
1/1/4            |        2          2 |        1          1 |        2   (Curr)
                 |        3          3 |        1          2 |        3   (Peak)
===============================================================================
```

Sample Output

```
A:PE-1# show subscriber-mgmt statistics sla-profiles
===============================================================================
SLA Profile Statistics
===============================================================================
SLA-Profile Name                        Current       Peak       Peak Timestamp
-------------------------------------------------------------------------------
tv-prof                                       0          0
hsi-prof                                      1          1  12/29/2016 15:50:50
hsi-iptv--prof                                0          0
no-prof-1                                     1          1  01/03/2017 20:00:37
hsi-voip-prof                                 0          2  12/29/2016 18:42:44
-------------------------------------------------------------------------------
Total                                         2
===============================================================================
```

Sample Output

```
A:PE-1# show subscriber-mgmt statistics sub-profiles
===============================================================================
Subscriber Profile Statistics
===============================================================================
Subscriber-Profile Name                  Current       Peak       Peak Timestamp
-------------------------------------------------------------------------------
no-prof                                       2          3  12/29/2016 18:42:44
sub-10M                                       0          0
sub-10M-UPNP                                  0          0
sub-20M                                       0          0
sub-20M-UPNP                                  0          0
-------------------------------------------------------------------------------
Total                                         2
===============================================================================
```

steering-profile

Syntax **steering-profile** *steering-profile-name*

Context show>subscr-mgmt

Description This command displays steering profile information.

Parameters *steering-profile-name* — Specifies an existing steering profile name.

sub-ident-policy

Syntax **sub-ident-policy** [*sub-ident-policy-name* [**association**]]
sub-ident-policy *sub-ident-policy-name* **script** {**primary** | **secondary** | **tertiary**}

Context show>subscr-mgmt

Description This command displays subscriber identification policy information.

Parameters *sub-ident-policy-name* — Specifies an existing subscriber identification policy name.

association — Displays information configured with the specified *sub-ident-policy-name*.

script {primary | secondary | tertiary} — Displays information for the specified identification script.

Output The following is an example of subscriber management subscriber identification policy information.

Sample Output

```
B:Dut-A>show>subscr-mgmt# sub-ident-policy
===============================================================================
Subscriber Identification Policies
===============================================================================
Name                            Description
-------------------------------------------------------------------------------
sub_ident_all
sub_ident_pc
-------------------------------------------------------------------------------
Number of Subscriber Identification Policies : 2
===============================================================================
B:Dut-A>show>subscr-mgmt#


B:Dut-A>show>subscr-mgmt# sub-ident-policy sub_ident_all
===============================================================================
Subscriber Identification Policy sub_ident_all
===============================================================================
Sub Profile Map
-------------------------------------------------------------------------------
Key                             Sub profile
```

```
-------------------------------------------------------------------------------
sub_prof100                     sub_prof100
sub_prof110                     sub_prof110
sub_prof120                     sub_prof120
sub_prof130                     sub_prof130
sub_prof140                     sub_prof140
sub_prof230                     sub_prof230
sub_prof80                      sub_prof80
sub_prof81                      sub_prof81
sub_prof90                      sub_prof90
-------------------------------------------------------------------------------
SLA Profile Map
-------------------------------------------------------------------------------
Key                             SLA profile
-------------------------------------------------------------------------------
sla_prof100_VOIP                sla_prof100_VOIP
sla_prof110_VOIP                sla_prof110_VOIP
sla_prof120_VOIP                sla_prof120_VOIP
sla_prof130_VOIP                sla_prof130_VOIP
sla_prof140_VOIP                sla_prof140_VOIP
sla_prof230_VOIP                sla_prof230_VOIP
sla_prof80_VOIP                 sla_prof80_VOIP
sla_prof81_VOIP                 sla_prof81_VOIP
sla_prof90_VOIP                 sla_prof90_VOIP
-------------------------------------------------------------------------------
Python Scripts
-------------------------------------------------------------------------------
#         Admin Oper  Script
          State State Name
-------------------------------------------------------------------------------
Primary   Down  Down  pyTom.py
Secondary Up    Up    pyTomDebug.py
Tertiary  Up    Up    hardcoded.py
===============================================================================
B:Dut-A>show>subscr-mgmt#
B:Dut-A>show>subscr-mgmt# sub-ident-policy sub_ident_all association
===============================================================================
Subscriber Identification Policy sub_ident_all
===============================================================================
SAP Associations
-------------------------------------------------------------------------------
Service-Id : 80 (VPLS)
 - SAP : 1/2/1:80
Service-Id : 90 (VPLS)
 - SAP : 1/2/1:90
Service-Id : 100 (VPLS)
 - SAP : 1/2/1:100
 - SAP : 1/2/1:101
 - SAP : 1/2/1:102
Service-Id : 110 (VPLS)
 - SAP : 1/2/1:110
 - SAP : 1/2/1:111
 - SAP : 1/2/1:112
Service-Id : 120 (VPLS)
 - SAP : 1/2/1:120
 - SAP : 1/2/1:121
 - SAP : 1/2/1:122
Service-Id : 130 (VPLS)
 - SAP : 1/2/1:130
```

```
Service-Id : 140 (VPLS)
 - SAP : 1/2/1:140
===============================================================================
B:Dut-A>show>subscr-mgmt#


B:Dut-A>show>subscr-mgmt# sub-ident-policy sub_ident_all script primary
===============================================================================
Subscriber Identification Policy sub_ident_all
===============================================================================
Primary Script
-------------------------------------------------------------------------------
URL         : ftp://xxx:yyy@a.b.c.d/pyTom.py
Admin State : Down                      Oper State : Down
-------------------------------------------------------------------------------
Source (dumped from memory)
-------------------------------------------------------------------------------
Script is not active.
-------------------------------------------------------------------------------
===============================================================================
B:Dut-A>show>subscr-mgmt#


B:Dut-A>show>subscr-mgmt# sub-ident-policy sub_ident_all script secondary
===============================================================================
Subscriber Identification Policy sub_ident_all
===============================================================================
Secondary Script
-------------------------------------------------------------------------------
URL         : ftp://xxx:yyy@a.b.c.d/pyTomDebug.py
Admin State : Up                        Oper State : Up
-------------------------------------------------------------------------------
Source (dumped from memory)
-------------------------------------------------------------------------------
      1 import alc
      2 yiaddr = alc.dhcp.yiaddr
      3 # Subscriber ID equals full client IP address.
      4 # Note: IP address 10.10.10.10 yields 'sub-168430090'
      5 # and not 'sub-10.10.10.10'
      6 alc.dhcp.sub_ident = 'sub-' + str(yiaddr)
      7 # DHCP server is configured such that the third byte (field) of the IP
      8 # address indicates the session Profile ID.
      9 alc.dhcp.sla_profile = 'sp-' + str((yiaddr & 0x0000FF00) >> 8)
===============================================================================
B:Dut-A>show>subscr-mgmt#


B:Dut-A>show>subscr-mgmt# sub-ident-policy sub_ident_all script tertiary
===============================================================================
Subscriber Identification Policy sub_ident_all
===============================================================================
Tertiary Script
-------------------------------------------------------------------------------
URL         : ftp://xxx:yyy@a.b.c.d/hardcoded.py
Admin State : Up                        Oper State : Up
-------------------------------------------------------------------------------
Source (dumped from memory)
-------------------------------------------------------------------------------
      1 from alc import dhcp
```

```
      2
      3 dhcp.sub_ident = 'sub_ident_A_1'
      4 dhcp.sub_profile_string = 'sub_prof_B_2'
      5 dhcp.sla_profile_string = 'sla_prof_C_3'
      6
===============================================================================
B:Dut-A>show>subscr-mgmt#
```

Table 145 describes subscriber identification policy output fields.

Table 145 **Subscriber Management Subscriber Identification Policy Field Descriptions**

Field	Description
Name	The subscriber identification policy name
Description	The user provided description of this policy
Number of Subscriber Identification Policies	The total number of subscriber identification policies configured on the node
Sub Profile Map	The subscriber profile used to define the aggregate QoS for all hosts within a subscriber context
Key	The key that defines the auto-generated subscriber identification key for PPP hosts.
Sub profile	A list of subscriber identification profiles
SLA profile	A list of SLA profiles
Admin State	The administrative state of the Python scripts
Oper State	The operational state of the Python scripts
Script Name	The Python policy script name
SAP Associations	The subscriber identification policy SAP associations
Service-Id	The service ID associated with the subscriber identification policy
SAP	The SAP ID associated with the subscriber identification policy
URL	The URL of the script to change the Python message
Admin State	The admin state of the script
Oper State	The operational state of the script
Primary Script	Information about the primary Python script
Secondary Script	Information about the secondary Python script

sub-profile

Syntax **sub-profile** [*sub-profile-name* [**association**]]

Context show>subscr-mgmt

Description This command displays subscriber profile information.

Parameters *sub-profile-name* — Specifies an existing subscriber profile name.

association — Displays the information configured with the specified *sub-profile-name*.

Output The following is an example of subscriber profile information.

Sample Output

```
A:Dut-A# show subscriber-mgmt sub-profile
===============================================================================
Subscriber Profiles
===============================================================================
Name                            Description
-------------------------------------------------------------------------------
sub_default
sub_prof100
sub_prof110
sub_prof120
sub_prof130
sub_prof140
sub_prof230
sub_prof80
sub_prof81
sub_prof90
sub_profPC1
sub_profPC2
sub_profPC3
-------------------------------------------------------------------------------
Number of Subscriber Profiles : 13
===============================================================================
A:Dut-A#


A:Dut-A# show subscriber-mgmt sub-profile sub_prof100
===============================================================================
Subscriber Profile sub_prof100
===============================================================================
I. Sched. Policy : service100
E. Sched. Policy : service100
Acct. Policy     : 1                                Collect Stats : Enabled
Last Mgmt Change : 07/10/2006 12:55:33
-------------------------------------------------------------------------------
Ingress Scheduler Overrides
-------------------------------------------------------------------------------
Scheduler                       Rate      CIR
-------------------------------------------------------------------------------
serv100                         8000      sum
-------------------------------------------------------------------------------
Egress Scheduler Overrides
```

```
-------------------------------------------------------------------------------
Scheduler                       Rate      CIR
-------------------------------------------------------------------------------
serv100                         8000      sum
-------------------------------------------------------------------------------
SLA Profile Map
-------------------------------------------------------------------------------
Key                             SLA Profile
-------------------------------------------------------------------------------
No mappings configured.
===============================================================================
A:Dut-A#


A:Dut-A# show subscriber-mgmt sub-profile sub_prof100 association
===============================================================================
Subscriber Profile sub_prof100
-------------------------------------------------------------------------------
SAP Default-Profile Associations
-------------------------------------------------------------------------------
No associations found.
-------------------------------------------------------------------------------
SAP Static Host Associations
-------------------------------------------------------------------------------
No associations found.
-------------------------------------------------------------------------------
SAP Non-Sub-Traffic-Profile Associations
-------------------------------------------------------------------------------
No associations found.
-------------------------------------------------------------------------------
Sub-Ident-Policy Profile Map Associations
-------------------------------------------------------------------------------
Policy-name : sub_ident_all
 - Key : sub_prof100
-------------------------------------------------------------------------------
Explicit Subscriber Map Associations
-------------------------------------------------------------------------------
No associations found.
===============================================================================
A:Dut-A#
```

Table 146

Table 146 **Subscriber Management Subscriber Profile Field Descriptions**

Field	Description
Name	The subscriber profile name
Description	The user provided description of this policy
Number of Subscriber Profiles	The total number of subscriber profiles configured on this node
Ingress Scheduler Overrides	Information about ingress scheduler overrides for the subscriber profile

Table 146 **Subscriber Management Subscriber Profile Field Descriptions**

Field	Description
Scheduler	The scheduler name
Rate	The scheduler rate
CIR	The CIR rate allowed for this scheduler
Egress Scheduler Overrides	Information about egress scheduler overrides for the subscriber profile
Key	The lookup key in the map
SLA Profile	Information about the SLA profile map
I. Sched. Policy	The ingress scheduler policy assigned to the ingress subscriber profile
E. Sched. Policy	The egress scheduler policy assigned to the egress subscriber profile
Acct. Policy	The number of accounting policies associated with this subscriber profile
Collect Stats	The accounting statistics collected on this subscriber profile.
Last Mgmt Change	The time of the most recent management-initiated change

accu-stats

Syntax **accu-stats subscriber** *subscriber-id*

Context show>subscriber-mgmt

Description This command displays the accumulated statistics for both online and offline statistics.

Parameters *subscriber-id* — Specifies the subscriber ID, up to 32 characters.

Output The following is an example of subscriber statistics.

Sample Output

```
A:active# show subscriber-mgmt accu-stats subscriber home-1
===============================================================================
Subscriber accumulated statistics for "home-1"
===============================================================================

-------------------------------------------------------------------------------
Subscriber accumulated statistics ingress queue 1
-------------------------------------------------------------------------------
```

```
                                    packets                octets
-------------------------------------------------------------------------------
Offered high priority               0                      0
Offered low priority                0                      0
Dropped high priority               0                      0
Dropped low priority                0                      0
Forwarded in-profile                0                      0
Forwarded out-of-profile            0                      0
-------------------------------------------------------------------------------

-------------------------------------------------------------------------------
Subscriber accumulated statistics egress queue 1
-------------------------------------------------------------------------------
                                    packets                octets
Forwarded in-profile                0                      0
Forwarded out-of-profile            0                      0
Dropped in-profile                  0                      0
Dropped out-of--profile             0                      0
-------------------------------------------------------------------------------
===============================================================================
```

Table 147 describes accumulated statistics output fields.

Table 147 **Subscriber Management Accumulated Statistics Field Descriptions**

Field	Description
packets	The number of packets associated with this policy
octets	The number of octets associated with this policy
Offered high priority	The number of offered high priority packets or octets
Offered low priority	The number of offered low priority packets or octets
Dropped high priority	The number of dropped high priority packets or octets
Forwarded in-profile	The number of forwarded in-profile packets or octets
Forwarded out-of-profile	The number of forwarded out-of-profile packets or octets
Dropped in-profile	The number of dropped in-profile packets or octets
Dropped out-of-profile	The number of dropped out-of-profile packets or octets

accu-stats-policy

Syntax **accu-stats-policy**
accu-stats-policy *policy-name* [**associations**]

Context show>subscriber-mgmt

Description This command displays the configuration of the accumulated statistics policy for the subscriber profile that referenced the accumulated statistics policy.

Parameters *policy-name* — Specifies the policy name, to a maximum of 32 characters.

associations — Lists the sub-profile that referenced the policy.

Output The following displays subscriber management accumulated statistics information.

Sample Output

```
A:active# show subscriber-mgmt accu-stats-policy "accu-stats-policy-01" associations
===============================================================================
Subscriber profile associations
===============================================================================
no-prof
sub-prof
-------------------------------------------------------------------------------
===============================================================================
No. of subscriber profiles: 2
===============================================================================

A:active# show subscriber-mgmt accu-stats-policy "accu-stats-policy-01"
===============================================================================
Subscriber accumulated statistics policy accu-stats-policy-01
===============================================================================
Description                         : (Not Specified)
Last management change              : 09/23/2016 16:48:46
===============================================================================

===============================================================================
Entries
===============================================================================
Direction Type     Id
-------------------------------------------------------------------------------
ingress   queue    1
egress    queue    1
===============================================================================
```

Table 148 describes subscriber management accumulated statistics policy fields.

Table 148 Subscriber Management Accumulated Statistics Policy Field Descriptions

Field	Description
Description	The user provided description of this policy
Last management change	The date and time of the last management change
Direction Type	The specific direction for the policer or the queue to be stored, ingress or egress
ID	The queue ID

Table 148 **Subscriber Management Accumulated Statistics Policy Field Descriptions (Continued)**

Field	Description
No. of subscriber profiles	The number of entries matching the search criteria

accu-stats-subscribers

Syntax **accu-stats-subscribers** [**subscriber** *sub-ident*] [**sub-profile** *name*]

Context show>subscriber-mgmt

Description This command lists all subscribers on the system that had referenced an accumulated statistics policy. The list contains subscribers that are currently referencing an **accu-stats-policy** as well as subscribers that no longer reference an **accu-stats-policy** (due to removal of the policy from the subscriber profile).

Parameters *sub-ident* — Specifies a specific subscriber ID, up to 32 characters.

sub-profile *name* — Lists all subscribers that are referencing the specified subscriber profile, up to 32 characters.

Output The following output displays subscriber management accumulated statistics information.

Sample Output

```
*DUT# show subscriber-mgmt accu-stats-subscribers
===============================================================================
Subscribers with accumulated statistics
===============================================================================
Subscriber                          : home-1
Subscriber profile                  : no-prof
Accumulated-stats policy            : accu-stats-policy-01
-------------------------------------------------------------------------------
Subscriber                          : home-2
Subscriber profile                  : no-prof
Accumulated-stats policy            : accu-stats-policy-02
-------------------------------------------------------------------------------
Subscriber                          : twamp-host
Subscriber profile                  : no-prof
Accumulated-stats policy            : (Unknown)
-------------------------------------------------------------------------------
No. of Subscribers with accumulated statistics: 3
===============================================================================
```

Sample Output

The **show subscriber-mgmt status system** command shows the number of subscribers using accumulated statistics and whether the statistics usage has reached the peak value.

```
*A:Dut-C# show subscriber-mgmt status system
===============================================================================
Subscriber Management System Status
===============================================================================
Chassis 1
-------------------------------------------------------------------------------
Memory usage high                                       : No
DHCP message processing overload                        : No
Statistics usage high                                   : Yes
Number of subscribers using statistics                  : 131072
Data-trigger statistics
-------------------------------------------------------------------------------
Packets received                                        : 0
Packets dropped                                         : 0
Packets in queue (actual)                               : 0
Packets in queue (peak)                                 : 0
Bridged Residential Gateway statistics
-------------------------------------------------------------------------------
BRG initialized                                         : 0
BRG operational                                         : 0
BRG in connectivity verification                        : 0
BRG on hold                                             : 0
BRG authenticated by proxy                              : 0
===============================================================================
```

Table 147 describes subscriber management status system fields.

Table 149 **Subscriber Management Status System Field Descriptions**

Field	Description
Chassis	The chassis number
Memory usage high	The memory usage of the subscriber management function on this chassis is near the limit; Yes or No
DHCP message processing overload	The subscriber management DHCP message processing function is in overload; Yes or No
Statistics usage high	The memory used for subscriber statistics on this chassis is near the limit; Yes or No
Number of subscribers using statistics	The number of subscribers using memory to store accumulated statistics
Packets received	The number of data packets received that may trigger creation of data-triggered subscriber hosts on this system
Packets dropped	The number of data packets dropped that could have triggered creation of data-triggered subscriber hosts on this system
Packets in queue (actual)	The number of data packets currently in the waiting queue that may trigger creation of data-triggered subscriber hosts on this system

Table 149 **Subscriber Management Status System Field Descriptions**

Field	Description
Packets in queue (peak)	The maximum value of the packets in queue object since the start up of this system or the last time that the value of packets in queue was reset
BRG initialized	The number of BRGs with the init timer running; waiting for the first host
BRG operational	The number of BRGs in the normal operating state, with dynamic hosts
BRG in connectivity verification	The number of BRGs with no hosts, pinging the residential gateway
BRG on hold	The number of BRGs with no hosts attached, in hold
BRG authenticated by proxy	The number of BRGs in this system

ancp-policy

Syntax **ancp-policy** [*policy-name*]

Context show>subscriber-management

Description This command displays subscriber ANCP policy information.

Output The following is an example of subscriber management ANCP policy information.

Sample Output

```
A:active# show subscriber-mgmt ancp-policy
===============================================================================
ANCP Policies
===============================================================================
adsl-operator1
vdsl-operator1
-------------------------------------------------------------------------------
Number of ANCP policies : 2
===============================================================================
A:active#

A:active# show subscriber-mgmt ancp-policy adsl-operator1
===============================================================================
ANCP Policy "adsl-operator1"
===============================================================================
I. Rate Reduction      : 0 kbps
I. Rate Adjustment     : 100 percent
I. Rate Monitor        : 0 kbps
I. Rate Monitor Alarm  : no
I. Rate Modify         : scheduler "root"
```

```
E. Rate Reduction      : 10 kbps
E. Rate Adjustment     : 100 percent
E. Rate Monitor        : 0 kbps
E. Rate Monitor Alarm  : no
E. Rate Modify         : scheduler "root"
Port Down : N/A
Last Mgmt Change: 01/26/2007 17:10:51
===============================================================================
A:active#

A:active# show subscriber-mgmt ancp-policy adsl-operator1 association
===============================================================================
ANCP Policy "adsl-operator1" associations
===============================================================================
SAP Static Map Associations
-------------------------------------------------------------------------------
- SAP     : 1/1/3                              Svc-id  : 333 (VPLS)
     String : "ANCP-String-1"
     String : "ANCP-String-2"
-------------------------------------------------------------------------------
MSS Static Map Associations
-------------------------------------------------------------------------------
- Cust-id : 1                                 MSS-name: mss1
     String : "ANCP-String-3"
-------------------------------------------------------------------------------
Subscriber Associations
-------------------------------------------------------------------------------
No associations found.
Number of associations : 3
===============================================================================
A:active#
```

Table 150 describes subscriber management ANCP policy fields.

Table 150 **Subscriber Management ANCP Policy Field Descriptions**

Field	Description
I. Rate Reduction	The ingress rate reduction applied for this policy
I. Rate Adjustment	The ingress rate adjustment percentage
I. Rate Monitor	The ingress rate monitor rate
I. Rate Monitor Alarm	The ingress rate monitor enabled; yes or no
I. Rate Modify	The ingress rate modify scheduler policy
E. Rate Reduction	The egress rate reduction applied for this policy percentage
E. Rate Adjustment	The egress rate adjustment
E. Rate Monitor	The egress rate monitor rate
E. Rate Monitor Alarm	The egress rate monitor enabled; yes or no

Table 150 **Subscriber Management ANCP Policy Field Descriptions**

Field	Description
E. Rate Modify	The egress rate monitor rate
Port Down	The actions taken on port-down
Last Mgmt Change	The time of the last modification of this entry
SAP	The ANCP static map entry SAP ID
Svc-id	The service ID associated with this static map
String	The ANCP string
Cust-id	The ID of the customer in which the multi-service-site specifies
MSS-name	The name of the multi-service-site used by this subscriber
Number of associations	The total number associations matching the search criteria

ancp-string

Syntax **ancp-string**
ancp-string *ancp-string*
ancp-string customer *customer-id* **site** *customer-site-name*
ancp-string sap *sap-id*

Context show>subscriber-management

Description This command displays subscriber ANCP string information.

Parameters *ancp-string* — Specify the ASCII representation of the DSLAM circuit-id name.

customer-id — Specify the associated existing customer name.

customer-site-name — Specify the associated customer's configured MSS name.

sap-id — Specifies the physical port identifier portion of the SAP definition.

Output The following output displays subscriber management ANCP string information.

Sample Output

```
A:active# show subscriber-mgmt ancp-string
===============================================================================
ANCP-Strings
===============================================================================
ANCP-String                                                 Assoc State
-------------------------------------------------------------------------------
"ANCP-String-1"                                             SAP   Up
"ANCP-String-2"                                             SAP   Down
```

```
"ANCP-String-3"                                                 MSS   Up
"ANCP-String-4"                                                 MSS   Unknown
"ANCP-String-5"                                                 ANCP  Up
"ANCP-String-6"                                                 MSS   Unknown
-------------------------------------------------------------------------------
Number of ANCP-Strings : 6
===============================================================================
A:active#


*A:Dut-C# show subscriber-mgmt ancp-string hpolSub43
===============================================================================
ANCP-String "hpolSub43"
===============================================================================
Type       : SUB - "hpolSub43"
State      : Up                 Ancp Policy: ancpPol
I. Rate    : 100 kbps           E. Rate    : 200 kbps
Adj I. Rate: N/A                Adj E. Rate: 200 kbps
Act I. Rate: N/A                Act E. Rate: 182 kbps
Service Id : 1 (VPRN)
Group      : Alu
Neighbor   : 10.100.100.1:49063
===============================================================================
*A:Dut-C#
```

Other applicable show command output:

```
A:active# show service id 333 sap 1/1/3 detail
===============================================================================
Service Access Points(SAP)
===============================================================================
Service Id         : 333
SAP                : 1/1/3                    Encap             : null
...
-------------------------------------------------------------------------------
ANCP Override
-------------------------------------------------------------------------------
Ing Sched Name: root
- PIR    : 100 kbps
- String : "ANCP-String-1"
Egr Sched Name: root
- PIR    : 100 kbps
- String : "ANCP-String-1"
-------------------------------------------------------------------------------
...
Dro. InProf           : 0                        0
Dro. OutProf          : 0                        0
===============================================================================
A:active#


A:active# show service customer 1 site mss1
===============================================================================
Customer  1
===============================================================================
Customer-ID        : 1
Description        : Default customer
...
-------------------------------------------------------------------------------
```

```
ANCP Override
-------------------------------------------------------------------------------
Egr Sched Name: root
- PIR    : 90 kbps
- String : "ANCP-String-3"
-------------------------------------------------------------------------------
Service Association
-------------------------------------------------------------------------------
No Service Association Found.
===============================================================================
A:active#
```

Table 151 describes ANCP string output fields.

Table 151 **Subscriber Management ANCP String Field Descriptions**

Field	Description
ANCP-String	The ANCP string name
Assoc	The entity associated with the ANCP string
State	The state of the entity associated with the ANCP string
Number of ANCP-Strings	The number of ANCP strings matching the search criteria
Type	The location the ANCP string is configured
Ancp Policy	The ANCP policy name
I. Rate	The ingress rate
E. Rate	The egress rate
Adj I. Rate	The adjusted ingress rate
Adj E. Rate	The adjusted egress rate
Act I. Rate	The actual ingress rate
Act E. Rate	The actual egress rate
Service Id	The service ID
Group	The group name
Neighbor	The neighbor address

radius-proxy-server

Syntax **radius-proxy-server** *server-name*
radius-proxy-server *server-name* **cache**
radius-proxy-server *server-name* **cache hex-key** *hex-string*

radius-proxy-server *server-name* **cache string-key** *string*
radius-proxy-server*server-name* **cache summary**
radius-proxy-server *server-name* **statistics**
radius-proxy-server

Context show>router

Description This command displays RADIUS proxy server information.

Parameters *server-name* — Specifies the default RADIUS proxy server name created in the **config>router>radius-proxy** context.

cache — Displays cached information.

hex-string — Displays the hex key string.

Values [0x0 to 0xFFFFFFFF (max 64 hex nibbles)]

string — Displays the packet type of the RADIUS messages to use to generate the key for the cache of this RADIUS proxy server.

summary — Displays summarized information.

statistics — Displays statistics for the specified RADIUS proxy server.

Output The following is an example of subscriber RADIUS proxy server

Sample Output

```
*B:asd-tr0610-dr421# show router radius-proxy-
server "ZiggoRadiusProxyAnyCast" statistics
===============================================================================
RADIUS Proxy server statistics for "ZiggoRadiusProxyAnyCast"
===============================================================================
Rx packet                                                 : 28454097
Rx Access-Request                                         : 24846521
Rx Accounting-Request                                     : 3607576
Rx dropped                                                : 22986
  Retransmit                                              : 22986
  Server admin down                                       : 0
  No RADIUS policy configured                             : 0
  No load-balance-key configured                          : 0
  Invalid length                                          : 0
  Invalid Code field                                      : 0
  Invalid attribute encoding                              : 0
  Invalid User-Name                                       : 0
  Invalid password                                        : 0
  Invalid accounting Authenticator                        : 0
  Invalid Message-Authenticator                           : 0
  No memory                                               : 0
  Accounting-Request with invalid Acct-Status-Type        : 0
  Accounting-Request with no Acct-Status-Type             : 0
  Registered user overload                                : 0
  Dropped by Python                                       : 0

Tx Access-Accept                                          : 1929725
```

```
Tx Access-Reject                                        : 302354
Tx Access-Challenge                                     : 22598950
Tx Accounting-Response                                  : 3598730
Tx dropped                                              : 1351
  No key to cache                                       : 0
  Cache key too long                                    : 0
  Cache attributes too long                             : 0
  Reached maximum number of cache entries               : 0
  No memory                                             : 0
  Server timeout                                        : 1351
  Server authentication failure                         : 0
  Server invalid Code                                   : 0
  Invalid attribute encoding                            : 0
  Registered user overload                              : 0
  No RADIUS server configured                           : 0
  RADIUS server send failure                            : 0
  Dropped by Python                                     : 0
===============================================================================


*B:asd-tr0610-dr421# show router radius-proxy-server "ZiggoRadiusDRPProxyanyCast-
LEG" statistics
===============================================================================
ISA RADIUS Proxy server statistics for "ZiggoRadiusDRPProxyanyCast-LEG"
===============================================================================
Group 1 member 1
-------------------------------------------------------------------------------
Rx packet                                               : 72250262
Rx Access-Request                                       : 61457394
Rx Accounting-Request                                   : 10792868
Rx dropped                                              : 1525690
  Retransmit                                            : 28470
  Wrong purpose                                         : 0
  No UE MAC to cache                                    : 1497212
  Client context limit reached                          : 0
  No ISA RADIUS policy configured                       : 0
  Invalid attribute encoding                            : 0
  Invalid password                                      : 0
  Accounting-Request with invalid Acct-Status-Type      : 0
  Accounting-Request with no Acct-Status-Type           : 0
  Invalid accounting Authenticator                      : 0
  Invalid Message-Authenticator                         : 8
  Management core overload                              : 0

Tx Access-Accept                                        : 5830313
Tx Access-Reject                                        : 743060
Tx Access-Challenge                                     : 54844862
Tx Accounting-Response                                  : 9294168
Tx dropped                                              : 12226
  Server timeout                                        : 12169
  Invalid response Authenticator                        : 57
  Invalid Message-Authenticator                         : 0
  Invalid attribute encoding                            : 0
  RADIUS server send failure                            : 0
```

Table 152 describes subscriber RADIUS proxy server output fields.

Table 152 **Subscriber RADIUS Proxy Server Field Descriptions**

Label	Description
Invalid response Authenticator Rx packet	The number of packets received by this RADIUS proxy server
Rx Access-Request	The number of Access-Request packets received by this RADIUS proxy server
Rx Accounting-Request	The number of Accounting-Request packets received by this RADIUS proxy server
Rx dropped	The number of packets received by this RADIUS proxy server but dropped
Retransmit	The number of packets received by this RADIUS proxy server that were rejected because they are retransmitted
Wrong purpose	The number of packets received by this RADIUS proxy server that were rejected because the value of tmnxRadProxSrvPurpose is set to a value not matching the type of packet
No UE MAC to cache	The number of packets received by this RADIUS proxy server that were rejected because the UE MAC address was not present in the packet
Client context limit reached	The number of packets received by this RADIUS proxy server that were rejected because the limit of client contexts was reached. For each RADIUS transaction a client context is created, and are deleted once the transaction is finished.
No ISA RADIUS policy configured	The number of packets received by this RADIUS proxy server that were rejected because it has no ISA RADIUS server policy configured for that type of packet
Server admin down	The number of packets received by this RADIUS proxy server that were rejected because it is administratively shut down
No RADIUS policy configured	The number of packets received by this RADIUS proxy server that were rejected because it has no RADIUS server policy configured for that type of packet
No load-balance-key configured	The number of packets received by this RADIUS proxy server that were rejected because the selected RADIUS server policy's algorithm is set to hashBased and no load balance key is configured
Invalid length	The number of packets received by this RADIUS proxy server that were rejected because their length was invalid

Table 152 **Subscriber RADIUS Proxy Server Field Descriptions (Continued)**

Label	Description (Continued)
Invalid Code field	The number of packets received by this RADIUS proxy server that were rejected because they had an invalid Code field
Invalid attribute encoding	The number of packets received by this RADIUS proxy server that were rejected because one of the attributes was incorrectly encoded
Invalid User-Name	The number of packets received by this RADIUS proxy server that were rejected because they contained an invalid User-Name attribute
Invalid password	The number of packets received by this RADIUS proxy server that were rejected because the User-Password attribute could not be decoded
Invalid accounting Authenticator	The number of accounting packets received by this RADIUS proxy server that were rejected because they contained an invalid Authenticator field
Invalid Message-Authenticator	The number of packets received by this RADIUS proxy server that were rejected because they contained an invalid Message-Authenticator attribute
Management core overload	The number of packets that were rejected by this RADIUS server because the ISA management core is not able to process any new RADIUS requests because of overload
No memory	The number of packets that were rejected by this RADIUS server because there was not enough memory to store them
Accounting-Request with invalid Acct-Status-Type	The number of accounting packets received by this RADIUS proxy server that were rejected because they contained an invalid Acct-Status-Type attribute
Accounting-Request with no Acct-Status-Type	The number of accounting packets received by this RADIUS proxy server that were rejected because they contained no Acct-Status-Type attribute
Registered user overload	The number of packets that were rejected by this RADIUS server because the registered user indicated to be in overload
Dropped by Python	The number of packets received by this RADIUS proxy server but dropped by Python
Tx Access-Accept	The number of Access-Accept packets transmitted by this RADIUS proxy server
Tx Access-Reject	The number of Access-Reject packets transmitted by this RADIUS proxy server

Table 152 **Subscriber RADIUS Proxy Server Field Descriptions (Continued)**

Label	Description (Continued)
Tx Access-Challenge	The number of Access-Challenge packets transmitted by this RADIUS proxy server
Tx Accounting-Response	The number of Accounting-Response packets transmitted by this RADIUS proxy server
Tx dropped	The number of packets dropped by this RADIUS proxy server before transmission
No key to cache	The number of packets that could not be cached by this RADIUS proxy server because the key information was not present in the packet
Cache key too long	The number of packets that could not be cached by this RADIUS proxy server because the key information present in the packet was too long
Cache attributes too long	The number of packets that could not be cached by this RADIUS proxy server because the total length of the attributes is too long
Reached maximum number of cache entries	The number of packets that could not be cached by this RADIUS proxy server because the limit has been reached
No memory	The number of packets that could not be transmitted by this RADIUS proxy server because there was not enough memory
Server timeout	The number of packets that were dropped because the RADIUS servers have timed out
Server authentication failure	The number of packets that were dropped because the RADIUS server replied with a packet which failed authentication (invalid response Authenticator or Message Authenticator attribute)
Server invalid Code	The number of packets that were dropped because the RADIUS server replied with a packet with an invalid Code field
Invalid attribute encoding	The number of packets that were dropped because the RADIUS server replied with a packet with an invalid attribute
Registered user overload	The number of packets that were dropped because the registered user indicated to be in overload
No RADIUS server configured	The number of packets that were dropped by this RADIUS server because the RADIUS server policy has no servers configured
RADIUS server send failure	The number of packets that were dropped by this RADIUS server because the packet could not get transmitted to one of the servers in the RADIUS server policy

Table 152 **Subscriber RADIUS Proxy Server Field Descriptions (Continued)**

Label	Description (Continued)
Dropped by Python	The number of packets that were dropped by this RADIUS server because the packet was dropped by the Python script
Invalid response Authenticator	The number of packets that were dropped because the RADIUS server replied with a packet which failed authentication

wpp

Syntax **wpp**
wpp [**portal** *wpp-portal-name*] [**host** *ip-address*] **hosts**
wpp portal *wpp-portal-name*
wpp statistics

Context show>router

Description This command displays WPP port-related information in the specified routing instance.

Parameters *wpp-portal-name* — Specifies the name of this WPP portal.

ip-address — Specifies the host IP address.

hosts — Displays the hosts enabled on the portal.

Output The following is an example of router WPP information.

Sample Output

```
show router wpp
===============================================================================
WPP portals
===============================================================================
Portal                          Address         Controlled-Rtr      Num-Itf
-------------------------------------------------------------------------------
svr1                            1.1.1.1         0                   0
svr2                            2.2.2.2         0                   0
-------------------------------------------------------------------------------
No. of portals: 2
===============================================================================


show router wpp portal "svr1"
===============================================================================
WPP Portal "svr1"
===============================================================================
Address                              : 1.1.1.1
Controlled router                    : 0
Number of enabled interfaces         : 0
Triggered hosts                      : disabled
Last management change               : 01/27/2014 00:48:45
```

```
===============================================================================
```

Table 153 describes router WPP output fields.

Table 153 **Router WPP Field Descriptions**

Portal	The web portal server name that the system communicates to for the hosts on the group interface
Address	The portal IP adress
Controlled-Rtr	The identifier of the virtual router instance being controlled by this portal
Num-Itf	The number of interfaces associated with this portal
Number of enabled interfaces	The number of interfaces associated with this portal that are administratively enabled
Triggered hosts	Triggered hosts are enabled or disabled on this portal
Last management change	The time of the most recent management-initiated change to this WPP portal

bonding

Syntax **bonding** [**detail**]
bonding bonding-id *bonding-id* [**detail**]

Context show>service

Description This command displays an overview of active bonding contexts and the associated access connections.

When a bonding ID is specified, only information for that specific context is displayed.

Parameters *bonding-id* — Specifies the ID of a specific bonding context.

detail — Displays additional detailed bonding settings, such as load-balancing.

Output The following is an example of bonding information.

Sample Output

```
Node# show service bonding
===============================================================================
Bonding Summary
===============================================================================
-------------------------------------------------------------------------------
Bond1 - svc:40 - itf:grp_itf_bonding
```

```
-------------------------------------------------------------------------------
 1 - PPP  - svc:30 - sap:1/1/3:1 - mac:00:00:00:00:00:01 - sid:1
 2 - IPOE - svc:11 - sap:pxc-1.b:1.8 - mac:00:03:ff:f0:01:00
-------------------------------------------------------------------------------
===============================================================================
show service bonding bonding-id Bond1 detail
===============================================================================
Bonding Summary
===============================================================================
-------------------------------------------------------------------------------
Bond1 - svc:40 - itf:grp_itf_bonding
-------------------------------------------------------------------------------
 1 - PPP  - svc:30 - sap:1/1/3:1 - mac:00:00:00:00:00:01 - sid:1
 2 - IPOE - svc:11 - sap:pxc-1.b:1.8 - mac:00:03:ff:f0:01:00
-------------------------------------------------------------------------------
Reference rate                : 10000000 kbps
Low threshold rate            : 8000000 kbps
High threshold rate           : 9000000 kbps
Initial load-balancing weight : 100%
Load-balancing weight change  : 5%
Current load-balancing weight : 100%
-------------------------------------------------------------------------------
===============================================================================
```

ipoe

Syntax **ipoe**

Context show>service>id

Description This command displays IPoE information.

gsmp

Syntax **gsmp**

Context show>service>id

Description This enables the command to display GSMP information.

neighbors

Syntax **neighbors group** [*name*] [*ip-address*]

Context show>service>id>gsmp

Description This command display GSMP neighbor information.

Parameters **group** — A GSMP group defines a set of GSMP neighbors which have the same properties.

name — Specifies a GSMP group name is unique only within the scope of the service in which it is defined.

ip-address — Specifies the ip-address of the neighbor.

Output The following output displays an example of service GSMP neighbor information.

Sample Output

These commands show the configured neighbors per service, regardless that there exists an open TCP connection with this neighbor. The admin state is shown because for a neighbor to be admin enabled, the service, gsmp node, group node and the neighbor node in this service must all be in the **no shutdown** state. Session gives the number of session (open TCP connections) for each configured neighbor.

```
A:active>show>service>id>gsmp# neighbors
===============================================================================
GSMP neighbors
===============================================================================
Group                          Neighbor           AdminState  Sessions
-------------------------------------------------------------------------------
dslam1                         192.168.1.2        Enabled     0
dslam1                         192.168.1.3        Enabled     0
-------------------------------------------------------------------------------
Number of neighbors shown: 2
===============================================================================
A:active>show>service>id>gsmp#

A:active>show>service>id>gsmp# neighbors group dslam1
===============================================================================
GSMP neighbors
===============================================================================
Group                          Neighbor           AdminState  Sessions
-------------------------------------------------------------------------------
dslam1                         192.168.1.2        Enabled     0
dslam1                         192.168.1.3        Enabled     0
-------------------------------------------------------------------------------
Number of neighbors shown: 2
===============================================================================
A:active>show>service>id>gsmp#


A:active>show>service>id>gsmp# neighbors group dslam1 192.168.1.2
===============================================================================
GSMP neighbors
===============================================================================
Group                          Neighbor           AdminState  Sessions
-------------------------------------------------------------------------------
dslam1                         192.168.1.2        Enabled     0
===============================================================================
A:active>show>service>id>gsmp#
```

The following table describes show service-id gsmp neighbors group output fields:

Table 154 **Show Service-ID GSMP Neighbors Group Field Descriptions**

Label	Description
Group	The GSMP group name
Neighbor	The neighbor IP address
AdminState	The administrative state of the neighbor
Sessions	The number of sessions (open TCP connections) for each configured neighbor

sessions

Syntax **sessions** [**group** *name*] **neighbor** *ip-address*] [**port** *port-number*] [**association**] [**statistics**]

Context show>service>id>gsmp

Description This command displays GSMP sessions information.

Parameters **group** — A GSMP group defines a set of GSMP neighbors which have the same properties.

name — Specifies a GSMP group name is unique only within the scope of the service in which it is defined.

ip-address — Specifies the ip-address of the neighbor.

port — Specifies the neighbor TCP port number use for this ANCP session.

Values 0 to 65535

association — Displays to what object the ANCP-string is associated.

statistics — Displays statistics information about an ANCP session known to the system.

Output The following output displays GSMP sessions information.

Sample Output

This show command gives information about the open TCP connections with DSLAMs.

```
A:active>show>service>id>gsmp# sessions
=======================================================
GSMP sessions for service 999 (VPRN)
=======================================================
Port   Ngbr-IpAddr     Gsmp-Group
-------------------------------------------------------
40590  192.168.1.2     dslam1
-------------------------------------------------------
Number of GSMP sessions : 1
=======================================================
A:active>show>service>id>gsmp#
```

```
A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590
===============================================================================
GSMP sessions for service 999 (VPRN), neighbor 192.168.1.2, Port 40590
===============================================================================
State           : Established
Peer Instance   : 1                  Sender Instance : a3cf58
Peer Port       : 0                  Sender Port     : 0
Peer Name       : 12:12:12:12:12:12  Sender Name     : 00:00:00:00:00:00
timeouts        : 0                  Max. Timeouts   : 3
Peer Timer      : 100                Sender Timer    : 100
Capabilities      : DTD OAM
Conf Capabilities : DTD OAM
Priority Marking  : dscp nc2
Local Addr.       : 192.168.1.4
Conf Local Addr.  : N/A
===============================================================================
A:active>show>service>id>gsmp#
A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590 association
===============================================================================
ANCP-Strings
===============================================================================
ANCP-String                                                  Assoc. State
-------------------------------------------------------------------------------
No ANCP-Strings found
===============================================================================
A:active>show>service>id>gsmp#
A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590 statistics
===============================================================
GSMP session stats, service 999 (VPRN), neighbor 192.168.1.2, Port 40590
===============================================================
Event                                  Received  Transmitted
---------------------------------------------------------------
Dropped                                0         0
Syn                                    1         1
Syn Ack                                1         1
Ack                                    14        14
Rst Ack                                0         0
Port Up                                0         0
Port Down                              0         0
OAM Loopback                           0         0
===============================================================
A:active>show>service>id>gsmp#
```

The following table describes service ID GSMP sessions output fields:

Table 155 **Show Service-ID GSMP Sessions Field Descriptions**

Label	Description
Port	The port ID number
Ngbr-IpAddr	The neighbor IP address
Gsmp-Group	The GSMP group ID

Table 155 **Show Service-ID GSMP Sessions Field Descriptions (Continued)**

Label	Description (Continued)
State	The GSMP state of this TCP connection
Peer Instance	The unique GSMP ID for each end of the GSMP connection
Peer Port	The unique GSMP ID for each end of the GSMP connection
Peer Name	The unique GSMP ID for each end of the GSMP connection
timeouts	The number of successive timeouts for this session
Peer Timer	The GSMP keepalive timer
Capabilities	The ANCP capabilities negotiated for this session
Conf Capabilities	The ANCP capabilities configured for this session
Priority Marking	The priority marking configured for this session
Local Addr	The IP address used by the box's side of the TCP connection
Conf. Local Addr.	The configured IP address used by the box's side of the TCP connection
Sender Instance	The instance sent to the neighbor in this session
Sender Port	The port sent to the neighbor in this session
Sender Name	The name sent to the neighbor in this session
Max. Timeouts	The maximum number of successive timeouts configured for this session
Sender Timer	The timeout value that is announced toward the neighbor. The neighbor uses this timeout value while waiting for an acknowledgment from this system.

Note: The association command gives an overview of each ANCP string received from this session.

```
A:active>show>service>id>gsmp# sessions neighbor 192.168.1.2 port 40590 association
===============================================================================
ANCP-Strings
===============================================================================
ANCP-String                                                    Assoc.
State
-------------------------------------------------------------------------------
7330-ISAM-E47 atm 1/1/01/01:19425.64048                        ANCP    Up
-------------------------------------------------------------------------------
Number of ANCP-Strings : 1
```

```
===============================================================================
A:active>show>service>id>gsmp
```

session

Syntax **session** [**sap** *sap-id*] [**mac** *ieee-address*] [**circuit-id** *circuit-id*] [**remote-id** *remote-id*] [**interface** *ip-int-name* | *ip-address*] [**inter-dest-id** *intermediate-destination-id*] [**no-inter-dest-id**] [**ip-address** *ip-prefix*[/*prefix-length*]] [**port** *port-id*] [**subscriber** *sub-ident-string*] [**sap-session-id** *sap-session-index*] [**detail**] [**wholesaler** *service-id*] [**router-advertisement-policy** *policy-name*]

Context show>service>id>ipoe

Description This command displays the identified IPoE session details active on the specified service instance.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

ieee-address — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

circuit-id — Specifies the circuit ID ,up to 127 characters.

remote-id — Specifies that information that goes into the remote-id sub-option in the DHCP relay packet, up to 255 characters.

ip-int-name — Specifies the name of the IP interface, up to 32 characters.

ip-address[/prefix-length] — Specifies information for the specified IP address and mask.

no-inter-dest-id — Displays the information about no intermediate destination ID.

port-id — Displays information about the specified port ID.

sap-session-index — Displays sap-session-index information.

service-id — Specifies the service ID of the wholesaler. This parameter applies only to the 7750 SR.

Values

service-id:	1 to 2147483647

intermediate-destination-id — Specifies the intermediate destination identifier which is encoded in the identification strings up to 32 characters.

sub-ident-string — Specifies an existing subscriber identification profile.

detail — Displays all IPoE session details.

policy-name — Specifies the name of the router advertisement policy, up to 32 characters.

Output The following is an example of service IPoE session information.

Sample Output

```
# show service id 4000 ipoe session
===============================================================================
IPoE sessions for svc-id 4000
===============================================================================
Sap Id                          Mac Address         Up Time         MC-Stdby
    Subscriber-Id
        [CircuitID] | [RemoteID]
-------------------------------------------------------------------------------
1/1/4:1201.27                   00:51:00:00:00:0c   0d 00:00:18
    ipoe-session-001
-------------------------------------------------------------------------------
CID | RID displayed when included in session-key
Number of sessions : 1
===============================================================================


# show service id 4000 ipoe session detail
===============================================================================
IPoE sessions for service 4000
===============================================================================
SAP                     : 1/1/4:1201.27
Mac Address             : 00:51:00:00:00:0c
Circuit-Id              : circuit-id-1
Remote-Id               : remote-id-1
Session Key             : sap-mac

MC-Standby              : No

Subscriber-interface    : sub-int-1
Group-interface         : group-int-1

Up Time                 : 0d 00:01:01
Session Time Left       : N/A
Last Auth Time          : 02/28/2015 01:01:09
Min Auth Intvl (left)   : 0d 00:05:00 (0d 00:03:59)
Persistence Key         : N/A

Subscriber              : "ipoe-session-001"
Sub-Profile-String      : "sub-profile-1"
SLA-Profile-String      : "sla-profile-1"
ANCP-String             : ""
Int-Dest-Id             : ""
App-Profile-String      : ""
Category-Map-Name       : ""
Acct-Session-Id         : "144DFF0000001354D806D5"
Sap-Session-Index       : 1

IP Address              : 10.10.1.201/24
IP Origin               : Radius
Primary DNS             : N/A
Secondary DNS           : N/A
Primary NBNS            : N/A
Secondary NBNS          : N/A
Address-Pool            : N/A

IPv6 Prefix             : 2001:db8:a:111::/64
IPv6 Prefix Origin      : Radius
```

```
IPv6 Prefix Pool        : ""
IPv6 Del.Pfx.           : 2001:db8:a001:a100::/56
IPv6 Del.Pfx. Origin    : Radius
IPv6 Del.Pfx. Pool      : ""
IPv6 Address            : 2001:db8:a:101::aaa:1
IPv6 Address Origin     : Radius
IPv6 Address Pool       : ""
Primary IPv6 DNS        : N/A
Secondary IPv6 DNS      : N/A

Radius Session-TO       : N/A
Radius Class            :
Radius User-Name        : 00:51:00:00:00:0c
-------------------------------------------------------------------------------
Number of sessions : 1
===============================================================================


*A:Dut-A# show service id 13 ipoe session router-advertisement-policy ra-policy-01
===============================================================================
IPoE sessions for svc-id 13
===============================================================================
Sap Id                          Mac Address         Up Time         MC-Stdby
    Subscriber-Id
        [CircuitID] | [RemoteID]
-------------------------------------------------------------------------------
1/1/1:13                        00:00:00:00:00:13   0d 00:13:06
    sub-1
1/1/1:13                        00:00:00:00:00:14   0d 00:08:03
    sub-1
1/1/1:13                        00:00:00:00:00:15   0d 00:06:23
    sub-1
-------------------------------------------------------------------------------
CID | RID displayed when included in session-key
Number of sessions : 3
===============================================================================
*A:Dut-A#
```

Table 156 describes IPoE session field information.

Table 156 **Service IPoE Session Field Descriptions**

Field	Description
SAP Id	The SAP ID of the IPoE session
Mac Address	The MAC address of the IPoE session
Up Time	The time of the last modification
MC-Stdby	The number of SAPs defined on this service on a port which is in multi-chassis standby mode
SAP	The SAP ID of the IPoE session
Circuit-Id	The circuit ID of the IPoE session

Table 156 **Service IPoE Session Field Descriptions (Continued)**

Field	Description
Remote-Id	The remote ID of the IPoE session
Session Key	The session key used for this IPoE session to group subscriber hosts
MC-Standby	The number of SAPs defined on this service on a port which is in multi-chassis standby mode
Subscriber-interface	The subscriber interface name
Group-interface	The group interface name
Up Time	The time of the last modification
Session Time Left	The time remaining for the session timeout
Last Auth Time	The date and time of the last authentication
Min Auth Intvl (left)	The minimum interval between two consecutive authentication attempts for the same ARP host
Persistence Key	The key value that is used to track this in the persistence file
Subscriber	The subscriber identification string
Sub-Profile-String	The subscriber profile name
SLA-Profile-String	The SLA profile name
ANCP-String	The value of the ancp-string received from either the DHCP or the RADIUS server
Int-Dest-Id	Intermediate destination identifier received from the RADIUS server
App-Profile-String	The application profile string
Category-Map-Name	The value of the category map name received from the RADIUS server
Acct-Session-Id	The accounting session ID
Sap-Session-Index	The sap session index of the PPP session
IP Address	The IP address
IP Origin	The origin of the source IP address used for connectivity verification
Primary DNS	The primary DNS server
Secondary DNS	The secondary DNS server

Table 156 **Service IPoE Session Field Descriptions (Continued)**

Field	Description
Primary NBNS	The primary NBNS server
Secondary NBNS	The secondary NBNS server
Address-Pool	The IP address pool used to search for an IP address for this session
IPv6 Prefix	The IPv6 prefix
IPv6 Prefix Origin	The origin of the IPv6 prefix
IPv6 Prefix Pool	The DHCP server pool from which the SLAAC prefix is allocated
IPv6 Del.Pfx.	The IPv6 delegated prefix
IPv6 Del.Pfx. Origin	The origin of the IPv6 delegated prefix
IPv6 Del.Pfx. Pool	The DHCP server pool from which the delegated prefix is allocated
IPv6 Address	The IPv6 address
IPv6 Address Origin	The origin of the IPv6 address
IPv6 Address Pool	The DHCP server pool from which the address is allocated
Primary IPv6 DNS	The primary DNS server
Secondary IPv6 DNS	The secondary DNS server
Radius Session-TO	The value to interpret the session timeout VSA from RADIUS
Radius Class	The Class attribute returned by the RADIUS server in an Access-Accept message
Radius User-Name	The User-Name attribute returned by the RADIUS server in an Access-Accept message
Number of session	The number of sessions returned from the search criteria

summary

Syntax **summary**

Context show>service>id>ipoe

Description This command displays a summary of IPoE information.

managed-hosts

Syntax **managed-hosts type** {**aaa** | **data-triggered**} [**ip** *ip-address*] [**mac** *ieee-address*]

Context show>service>id

Description This command displays information about hosts that are not yet linked to an in-band control plane protocol, such as DHCP or DHCPv6. Examples include hosts installed by data-triggers and hosts installed by RADIUS.

Parameters **aaa** — Displays information about managed hosts installed and removed only by RADIUS.

data-triggered — Displays information about managed hosts installed by data-triggers.

ip-address — Displays information about a managed host with the specified IPv4 or IPv6 address.

ieee-address — Displays information about a managed host with the specified MAC address.

Output The following is an example of managed host information.

Sample Output

```
# show service id 100 managed-hosts type data-triggered
===============================================================================
Managed data-triggered hosts
===============================================================================
IP address                                MAC address
-------------------------------------------------------------------------------
10.0.0.9/32                               fa:ac:a6:02:11:01
2001:a:b:5::1/128                         fa:ac:a6:02:11:01
-------------------------------------------------------------------------------
No. of Managed hosts: 2
===============================================================================
```

slaac

Syntax **slaac**

Context show>service>id

Description This command displays SLAAC information.

host

Syntax **host** [**detail**] [**wholesaler** *service-id*] [**session** {**none** | **ipoe** | **ppp**}] [**router-advertisement-policy** *policy-name*]

host interface *interface-name* [**detail**] [**wholesaler** *service-id*] [**session** {**none** | **ipoe** | **ppp**}]

[**router-advertisement-policy** *policy-name*]
host mac *ieee-address* [**detail**] [**wholesaler** *service-id*] [**session** {**none** | **ipoe** | **ppp**}] [**router-advertisement-policy** *policy-name*]
host ipv6-address *ipv6-prefix* [**detail**] [**wholesaler** *service-id*] [**session** {**none** | **ipoe** | **ppp**}] [**router-advertisement-policy** *policy-name*]
host sap *sap-id* [**detail**] [**wholesaler** *service-id*] [**session** {**none** | **ipoe** | **ppp**}] [**router-advertisement-policy** *policy-name*]

Context show>service>id>slaac

Description This command displays SLAAC host related information.

Parameters **detail** — Specifies to display detailed SLAAC host related information.

service-id — Specifies the service ID of the wholesaler.

Values 1 to 2147483647

session — Specifies to display information about SLAAC hosts that are associated with an IPoE session or a PPP session.

Values none, ipoe, ppp

interface-name — Specifies the interface name, up to 32 characters.

ieee-address — Specifies the MAC address for which to display SLAAC host information.

ipv6-address — Specifies the IPv6 address for which to display SLAAC host information.

sap-id — Specifies the physical port identifier portion of the SAP definition.

policy-name — Specifies the name of the router advertisement policy.

Output The following is an example of detailed SLAAC host information.

Sample Output

```
*A:eng-BNG-2# show service id 1000 slaac host detail
===============================================================================
SLAAC hosts for service 1000
===============================================================================
Service ID           : 1000
Prefix               : 2001:1000:bad:beef::/64
Interface Id         : N/A
Mac Address          : 00:00:10:10:12:12
Subscriber-interface : sub-int-01
Group-interface      : grp-int-01
SAP                  : 1/1/20:841
Termination Type     : local
Creation Time        : 2017/02/13 22:05:03
Persistence Key      : N/A
Router adv. policy   : ra-policy-01
IPoE|PPP session     : No
Radius sub-if prefix : N/A
IPoE Trigger         : rtr-solicit
Last Auth Time       : 2017/02/13 22:05:03
Inactivity Timer     : N/A
Sub-Ident            : "host-1"
```

```
Sub-Profile-String   : "no-prof"
SLA-Profile-String   : "no-prof"
App-Profile-String   : ""
ANCP-String          : ""
Int Dest Id          : ""
Category-Map-Name    : ""
Info origin          : radius
Pool                 : ""
Primary-Dns          : N/A
Secondary-Dns        : N/A
Circuit Id           : N/A
Remote Id            : N/A
-------------------------------------------------------------------------------
Number of hosts : 1
===============================================================================

*A:Dut-A# show service id 13 slaac host router-advertisement-policy ra-policy-01
===============================================================================
SLAAC hosts for service 13
===============================================================================
Prefix                     Mac Address          Sap Id                MC-Stdby
  IntId
-------------------------------------------------------------------------------
2013:bad:beef:1::/64       00:00:00:00:00:13    1/1/1:13
  N/A
2013:bad:beef:2::/64       00:00:00:00:00:14    1/1/1:13
  N/A
2013:bad:beef:3::/64       00:00:00:00:00:15    1/1/1:13
  N/A
-------------------------------------------------------------------------------
Number of hosts : 3
===============================================================================
*A:Dut-A#
```

Table 157 describes service SLAAC host information.

Table 157 **Service SLAAC Host Field Descriptions**

Field	Description
Service ID	The service identifier.
Prefix	The prefix for the service
Inteface Id	The interface identifier of this SLAAC host. If the SLAAC host has multiple addresses, this field will only show the interface identifier of one of them.
Mac Address	The MAC address of the this SLAAC hos
Subscriber-interface	The subscriber interface name
Group-interface	The subscriber interface group interface name
SAP	The service SAP ID

Table 157 **Service SLAAC Host Field Descriptions (Continued)**

Field	Description
Termination Type	The termination type of this SLAAC host
Creation Time	The date and time the service was created
Persistence Key	The key value that can be used to track this SLAAC host in the persistence file.
Router adv. policy	The name of the router advertisement policy associated with the SLAAC host during authentication
IPoE|PPP session	The indication if this SLAAC host belongs to an IPoE or PPP session
Radius sub-if prefix	The subscriber interface prefix that is received from a Radius server
IPoE Trigger	The trigger name
Last Auth Time	The sysUpTime at the time of the last authentication for this session
Inactivity Timer	The remaining time in seconds before the host is removed. A value of 0 indicates no inactivity timer is active.
Sub-Ident	The name of the subscriber identification policy
Sub-Profile-String	The name of the subscriber profile
SLA-Profile-String	The name of the SLA profile
App-Profile-String	The application profile string
ANCP-String	The ANCP string applicable for this session
Int Dest ID	The intermediate destination identifier applicable for this session
Category-Map-Name	The category map name applicable for this session
Info origin	The server that provided the Sub-profile, SLA profile, Ancp string and Intermediate Destination Id
Pool	the DHCP server pool from which the SLAAC prefix is allocated
Primary-Dns	the primary DNS server
Secondary-Dns	the secondary DNS server
Circuit ID	circuit ID of this host
Remote ID	remote ID of this host

Table 157 **Service SLAAC Host Field Descriptions (Continued)**

Field	Description
Number of hosts	The total number of hosts matching the search criteria
MC-Stdby	The number of SAPs defined on this service on a port which is in multi-chassis standby mode

summary

Syntax **summary**

Context show>service>id>slaac

Description This command displays a summary of all SLAAC hosts.

subscriber-hosts

Syntax **subscriber-hosts** [**sap** *sap-id*] [**ip** *ip-prefix/prefix-length*] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**app-profile** *app-profile-name*] [**wholesaler** *service-id*] [**address-origin** *address-origin*] [**detail**] [**statistics**] [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]

Context show>service>id

Description This command displays subscriber host information.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

ip-prefix/prefix-length — Shows information for the specified IP address and mask.

Values *ipv4-prefix* — a.b.c.d (host bits must be 0)
ipv4-prefix-length — 0 to 32
ipv6-prefix — x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x — [0..FFFF]H
d — [0..255]D
ipv6-prefix-length — 0 to 128

ieee-address — Displays information only for the specified 48-bit MAC address. The MAC address can be expressed in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.

sub-profile-name — Displays an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config**>**subscr-mgmt**>**sub-profile** context.

sla-profile-name — Displays subscriber host information for hosts with the specified SLA profile name.

app-profile-name — Specifies an existing SLA app profile name, up to 32 characters.

service-id — Specifies the ID that uniquely identifies a service.

Values service-id: 1 to 214748364
svc-name: string, up to 64 characters.

address-origin — Specifies an address origin.

Values aaa, dynamic, static, bonding

detail — Displays detailed information.

statistics — Displays the HTTP redirect statistics for the host. The redirect statistics are divided into HTTP GET and HTTP response for redirection.

spi-sharing-type — Displays subscriber host information for hosts with the specified SPI sharing type.

Values per-group, per-session-ipoe, per-session-ppp

spi-sharing-id — Displays subscriber host information for hosts with the specified SPI sharing type and sharing identifier. The SPI group ID is:

- a sharing ID when the sharing type is per-group
- an internal PPP or IPoE sharing ID when the sharing type is per-PPP or per-IPoE session

Values 0 to 4294967295

Output The following output displays an example of subscriber host information.

Sample Output

```
*A:Dut-C>#  show service id 3 subscriber-hosts
===============================================================
Subscriber Host table
===============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID     Origin       Fwding State
---------------------------------------------------------------
2/1/5:2                    TEACAHEH74
  10.11.1.61
    00:80:00:00:00:0a          N/A               ARP-Host        Fwding
[pw-11:11]                  VIACAHEH74
  10.11.1.2
    00:00:11:11:01:02          N/A               ARP-Host        Fwding
[pw-11:12]                  pw-11:12
  10.11.1.3
    00:00:11:11:01:03          N/A               ARP-Host        Fwding
[pw-11:13]                  pw-11:13
  10.11.1.4
    00:00:11:11:01:04          N/A               ARP-Host        Fwding
[pw-22:22]                  XMACAHEH74
  10.22.1.2
```

```
    00:00:22:22:01:02           N/A              ARP-Host        Fwding

[pw-33:33]                   IUASAHEH74
  10.33.1.2
    00:00:33:33:01:02           N/A              ARP-Host        Fwding
-------------------------------------------------------------
Number of subscriber hosts: 6
=============================================================
*A:Dut-C>#


A:Dut-A# show service id 100 subscriber-hosts ip 10.100.1.5
===============================================================================
Subscriber Host table
===============================================================================
Sap                    IP Address      MAC Address       Origin(*) Subscriber
-------------------------------------------------------------------------------
1/2/1:102              10.100.1.5      00:10:00:00:00:03 -/D/- nokia_100
-------------------------------------------------------------------------------
Number of subscriber hosts : 1
===============================================================================
(*) S=Static Host, D=DHCP Lease, N=Non-Sub-Traffic
A:Dut-A#
```

The following table describes **show service-id subscriber hosts** output fields:

Table 158 **Show Service-ID Subscriber Hosts Field Descriptions**

Label	Description
Sap	The physical port identifier portion of the SAP definition.
IP Address	The IP address
MAC Address	The MAC address
PPPoE-SID	The PPPoE-SID of the originating subscriber
Origin Subscriber	The ID of the originating subscriber
Fwding State	The forwarding state
Number of subscriber hosts	The number of subscriber hosts

subscriber-using

Syntax **subscriber-using** [**service-id** *service-id*] [**sap-id** *sap-id*] [**interface** *ip-int-name*] [**ip** *ip-address*[/*mask*]] [**mac** *ieee-address*] [**sub-profile** *sub-profile-name*] [**sla-profile** *sla-profile-name*] [**app-profile** *app-profile-name*] [**port** *port-id*] [**no-inter-dest-id** | **inter-dest-id** *intermediate-destination-id*]

Context show>service

Description This command displays selective subscriber information using specific options.

Parameters *service-id* — Specifies the ID that uniquely identifies a service.

sap-id — Displays the physical port identifier portion of the SAP definition.

ip-int-name — Shows DHCP statistics on the specified interface.

port-id — Indicates the SAP or SDP for which this entry contains information.

ip-address[/*mask*] — Shows information for the specified IP address and mask.

ieee-address — Displays information only for the specified 48-bit MAC address. The MAC address can be expressed in the form *aa*:*bb*:*cc*:*dd*:*ee*:*ff* or *aa-bb-cc-dd-ee-ff* where *aa*, *bb*, *cc*, *dd*, *ee* and *ff* are hexadecimal numbers.

sub-profile-name — Displays an existing subscriber profile name to be associated with the static subscriber host. The subscriber profile is configured in the **config>subscr-mgmt>sub-profile** context.

sla-profile-name — Specifies an existing SLA profile name.

app-profile-name — Displays the profile of the specified application.

intermediate-destination-id — Indicates the intermediate destination identifier received from either the DHCP or RADIUS server, or the local user database.

Output The following output displays an example of selective subscriber information.

Sample Output

```
A:Dut-A# show service subscriber-using service-id 100
===============================================================================
Subscribers
===============================================================================
Subscriber                        Sub Profile
-------------------------------------------------------------------------------
nokia_100                      sub_prof100
-------------------------------------------------------------------------------
Matching Subscribers : 1
===============================================================================
A:Dut-A#

A:Dut-A# show service subscriber-using
===============================================================================
Subscribers
===============================================================================
Subscriber                        Sub Profile
-------------------------------------------------------------------------------
nokia_100                         sub_prof100
nokia_110                         sub_prof110
nokia_120                         sub_prof120
nokia_130                         sub_prof130
nokia_80                          sub_prof80
nokia_90                          sub_prof90
client_PC1                        sub_profPC1
static                            sub_default
-------------------------------------------------------------------------------
Matching Subscribers : 8
```

```
===============================================================
A:Dut-A#
```

Table 159 describes subscriber using a specified service ID field descriptions.

Table 159 **Subscriber Using Service ID Field Descriptions**

Field	Description
Subscriber	The subscriber name
Sub Profile	The subscriber profile name
Matching Subscribers	The number of subscribers matching the search criteria

call-trace

Syntax **call-trace**

Context show

Description This command enables the context to display information related to the call-trace module.

files

Syntax **files** [**running** | **finished**] [**detail**]

Context show>call-trace

Description This command gives an overview of all the files in use by the call-trace module, either for running or finished jobs.

Parameters **running** — Displays the active jobs that are tracing events triggered by the host being monitored.

finished — Displays the jobs that have already finished and not tracing events generated by the host anymore.

detail — Displays detailed information about the files.

Output The following output displays call trace local log files information.

Sample Output

```
Node# show call-trace files
===============================================================================
Call trace files of running jobs
===============================================================================
File path                                                          File size
                                                                     (bytes)
```

```
-------------------------------------------------------------------------------
cf1-A:/calltrace/running/mac_00-02-00-00-00-19_160119_1557.pcap            1575
-------------------------------------------------------------------------------
No. of call trace files of running jobs: 1
===============================================================================
===============================================================================
Call trace files of finished jobs
===============================================================================
File path                                                            File size
                                                                       (bytes)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
No. of call trace files of finished jobs: 0
===============================================================================
```

Table 160 describes call trace output fields.

Table 160 **Call Trace Field Descriptions**

Field	Description
File path	The compact flash (CF) configuration to store call trace files.
File size	The total size of call-trace files on the specified compact flash card
No. of call trace files running jobs	The total number of active trace call files running jobs
No. of call trace files of finished jobs	The total number of finished trace call files jobs

ipoe

Syntax **ipoe**

Context show>call-trace

Description This command enables the context to display IP over Ethernet (IPoE) call trace information.

session

Syntax **session** [**name** *trace-name*] [**detail**]
session sap *sap-id* **mac** *ieee-address* [{**circuit-id** *circuit-id* | **remote-id** *remote-id*}] [**detail**]

Context show>call-trace>ipoe

Description This command displays call trace information for the active IPoE session.

Configuring this command with a specific IPoE session key displays information for the specified trace job only.

Parameters *circuit-id* — Specifies a circuit ID that identifies an IPoE session.

detail — Displays detailed information.

ieee-address — Specifies a MAC address that identifies an IPoE session.

remote-id — Specifies a remote ID that identifies an IPoE session.

trace-name — Specifies the name of the trace that triggered the job.

sap-id — Specifies the SAP ID that identifies an IPoE session.

Output The following output displays call trace information for the active IPoE session.

Sample Output

```
Node # show call-trace ipoe session
=========================================================
Call trace IPoE session
=========================================================
IEEE address             Status          No. of messages
  Trace name
---------------------------------------------------------
00:02:00:00:00:19        running                     10
     default
---------------------------------------------------------
No. of call trace IPoE debug jobs: 1
=========================================================


Node# show call-trace ipoe session detail
===============================================================================
Call trace IPoE session
===============================================================================
SAP ID      : 4/1/ 2
IEEE address: 00:02:00:00:00:19
Trace name  : default                           Status       : running
Size limit  : 10 MB                             Time limit   : 86400 sec
Live output
-------------------------------------------------------------------------------
Destination : none
Statistics
-------------------------------------------------------------------------------
Started               : 2017/01/11 15:34:08UTC
No. of captured msgs  : 10
Size of captured msgs : 2575 B
===============================================================================
```

Table 161 describes call trace IPoE session output fields.

Table 161 **Call Trace IPoE Session Field Descriptions**

Field	Description
IEEE address	The IEEE address of the equipment to match by this trace
Status	The status of the trace session
No. of messages	The number of messages from the trace session
Trace name	The name of the trace profile
Size limit	The maximum data volume generated by a single call trace job to the output, in megabytes
Time limit	The maximum duration of a single call trace job, in seconds
Live output	The live output destination for trace logs
Destination	The destination address of the trace logs
Started	The date and time the trace began
No. of captured msgs	The number of captured messages
Size of captured msgs	The size of the captured messages

trace

Syntax **trace** [**name** *trace-name*]

Context show>call-trace>ipoe

Description This command displays information about all traces and associated parameters that are currently enabled on the system. Including the *trace-name* parameter limits the output to the specified trace only.

Parameters *trace-name* — Specifies the name of the trace to be displayed.

Output The following output displays call trace information.

Sample Output

```
Node# show call-trace ipoe trace
===============================================================================
Call trace trace
===============================================================================
Trace name  : default                           Type         : ipoe
SAP ID      :
IEEE address: 00:02:00:00:00:19
Maximum jobs: 1
===============================================================================
```

Table 162 describes call trace IPoE trace output fields.

Table 162 **Call Trace IPoE Trace Field Descriptions**

Field	Description
Trace name	The name of the trace profile
Type	The type of the call trace
SAP ID	The SAP ID of the session to match by this trace
IEEE address	The IEEE address of the session that matches this trace
Maximum jobs	The maximum number of call trace jobs that are started by this trace

status

Syntax **status**

Context show>call-trace

Description This command gives a router-wide overview of call-trace operational data, such as number of configured profile, number of jobs and status of the compact flash.

Output The following output displays call trace status information.

Sample Output

```
===============================================================================
Call trace status
===============================================================================
No. of trace profiles: 1
No. of trace jobs    : 0
-------------------------------------------------------------------------------
Primary CF           : CF1              Max files number     : 200
CF1 state            : enabled          CF1 volume limit     : 1000 MB
CF2 state            : enabled          CF2 volume limit     : 1000 MB
```

Table 163 describes call trace status output field information.

Table 163 **Call Trace Status Field Descriptions**

Field	Descriptions
No. of trace profiles	The number of profiles configured in the system
No. of trace jobs	The number of trace jobs

Table 163 **Call Trace Status Field Descriptions (Continued)**

Field	Descriptions
Primary CF	The compact flash card to be used as the primary local storage location to save the generated call trace log files
Max files number	The cumulative number of call trace log files present on all cflash cards on the active CPM that are being used for local storage has reached the limit
CF1 state	The status of a call trace job on CF2. “running” indicates the job is active and is tracing events triggered by the host being monitored. “finished” indicates the job has already finished and is not tracing events generated by the host anymore.
CF1 volume limit	The data volume limit for the CF1 call trace job
CF2 state	The status of a call trace job on CF2. “running” indicates the job is active and is tracing events triggered by the host being monitored. “finished” indicates the job has already finished and is not tracing events generated by the host anymore.
CF2 volume limit	The data volume limit for the CF2 call trace job

trace-profile

Syntax **trace-profile** [**detail**]
trace-profile *profile-name*

Context show>call-trace

Description This command provides an overview of all configured profiles or details of a specific profile. If the detail option is specified the full information for all configured profiles are displayed.

Parameters **detail** — Displays detailed trace profile information.

Output The following output displays call-trace trace-profile information.

Sample Output

```
Node# show call-trace trace-profile "default"
===============================================================================
Call-trace  trace profile
===============================================================================
Profile name              : default
Description               : none
Live output               : none
Format                    : pcap
Size limit                : 10 MB
```

```
Time limit               : 86400 secs
-------------------------------------------------------------------------------
Number of profiles : 1
===============================================================================
```

dhcp-python-policy

Syntax **dhcp-python-policy** *policy-name*
no dhcp-python-policy

Context config>service>vpls>sap

Description This command specified the Python policy for DHCPv4 packets sent/received on the capture SAP.

The **no** form of this command removes the policy name from the configuration.

Parameters *policy name* — Specifies an existing Python policy name

dhcp-user-db

Syntax **dhcp-user-db** *local-user-db-name*
no dhcp-user-db

Context config>service>vpls>sap

Description This command enabled access to LUDB for DHCPv4 hosts under the capture SAP. The name of this local user database must match the name of local user database configured under the **config>service>vprn/ies>sub-if>group-if>dhcp** context.

Parameters *local-user-db* — Specifies the name of the local user database name up to 32 characters.

dhcp6-python-policy

Syntax **dhcp6-python-policy** *policy-name*
no dhcp6-python-policy

Context config>service>vpls>sap

Description This command specified the Python policy for DHCPv6 packets sent/received on the capture SAP.

The **no** form of this command removes the policy name from the configuration.

Parameters *policy name* — Specifies an existing Python policy name

dhcp6-user-db

Syntax **dhcp6-user-db** *local-user-db*
no dhcp6-user-db

Context config>service>vpls>sap

Description This command enabled access to LUDB for DHCPv6 hosts under the capture SAP. The name of this LUDB must match the name of the LUDB configured under the **config>service>vprn/ies>sub-if>grp-if>dhcp** context.

The **no** form of this command reverts to the default.

Parameters *local-user-db* — Specifies the name of the local user database name up to 32 characters.

9.25.2.3 Clear Commands

session

Syntax **session sap** *sap-id* **mac** *ieee-address* [{**circuit-id** *circuit-id* | **remote-id** *remote-id*}]

Context clear>call-trace>ipoe

Description This command clears the trace job identified by this session.

The *sap*, *mac*, *circuit-id*, and *remote-id* parameters cannot use wildcard characters, and must match the IPoE session key. This command does not affect any of the debug trace commands, and new jobs can still be executed for other IPoE sessions.

After a session is cleared, tracing for the session will not be restarted by any configured trace. Only explicitly starting a new trace with the **trace-existing-sessions** parameter can restart tracing for the session.

Parameters *circuit-id* — Specifies a circuit ID that identifies an IPoE session.

ieee-address — Specifies a MAC address that identifies an IPoE session.

remote-id — Specifies a remote ID that identifies an IPoE session

sap-id — Specifies the SAP ID that identifies an IPoE session.

trace

Syntax **trace** *trace-name*

Context clear>call-trace>ipoe

Description This command clears all trace jobs started by the specified trace. This command does not affect the trace command itself, and new jobs can still be executed for other IPoE sessions.

After a session is cleared, tracing for the session will not be restarted by any configured trace. Only explicitly starting a new trace with the **trace-existing-sessions** parameter can restart tracing for the session.

Parameters *trace-name* — Specifies the configured name of the trace.

accu-stats

Syntax **accu-stats active-subs no-accu-stats-policy**
accu-stats active-subs sub-profile *profile-name*
accu-stats inactive-subs
accu-stats subscriber *subscriber-id*

Context clear>subscriber-mgmt

Description This command removes the accumulated statistics for a subscriber.

Parameters **active-subs no accu-stats-policy** — Removes the accumulated statistics that are no longer associated with an **accu-stats-policy**. These are active subscribers that are referencing a subscriber profile that removed an **accu-stats-policy**.

active-subs sub-profile *profile-name* — Removes the accumulated statistics for the specified subscriber profile.

inactive-subs — Removes the accumulated statistics from all inactive subscribers.

profile-name — Specifies the subscriber profile up to 32 characters.

subscriber-id — Specifies the subscriber up to 32 characters.

ancp-sub-string

Syntax **ancp-sub-string** *string*

Context clear>subscr-mgmt>ancp>ancp

Description This command clears subscriber ANCP data.

Parameters *string* — Clears the ANCP string corresponding to this subscriber ID.

arp

Syntax **arp** {**all** | *ip-address*}
arp interface [*ip-int-name* | *ip-address*]

Context clear>router

Description This command clears all or specific ARP entries.

The scope of ARP cache entries cleared depends on the command line option specified.

Parameters **all** — Clears all ARP cache entries

ip-addr — Clears the ARP cache entry for the specified IP address.

ip-int-name — Clears all ARP cache entries for the interface with the specified name.

ip-addr — Clears all ARP cache entries for the specified interface with the specified address.

authentication

Syntax **authentication** [*policy-name*]
authentication coa-statistics

Context clear

Description This command clears subscriber authentication data.

Parameters *policy-name* — Clears the specifies existing authentication policy name.

coa-statistics — Clears statistics for incoming RADIUS Change of Authorization requests.

bonding

Syntax **bonding** *bonding-id*

Context clear>subscr-mgmt

Description This command clears the bonding subscriber context and all related access sessions.

Parameters *bonding-id* — Specifies the identifier of the bonding context as passed in AAA to the access host.

diameter-session

Syntax **diameter-session**

Context clear>subscr-mgmt

Description This command clears diameter session data.

ccrt-replay

Syntax **ccrt-replay diameter-application-policy** *name* **statistics**
ccrt-replay diameter-application-policy *name* **sessions**

Context clear>subscr-mgmt>diam-session

Description This command, when the **statistics** parameter is specified, clears the CCR-t replay statistics and updates the “Statistics last cleared time” value.

This command, when the **sessions** parameter is specified, drops the diameter Gx or Gy sessions that are in CCR-t replay mode.

Parameters *name* — Specifies the diameter application policy, up to 32 characters.

statistics — Clears the CCR-t replay statistics.

sessions — Clears the sessions that are in CCR-t replay mode.

errors

Syntax **errors**

Context clear>subscr-mgmt

Description This command clears all errors in the circular buffer.

idle-only-msap

Syntax **idle-only-msap** *msap-policy-name*

Context clear>subscr-mgmt

Description This command removes all idle MSAPs associated with the MSAP policy.

This command only removes idle MSAPs without active subscribers. This command is considered safer than the **clear**>**subscr-mgmt**>**msap-policy** *msap-policy-name* **idle-only** command because in that command, the parameter **idle-only** is optional. Not specifying **idle-only** deletes MSAPs with active subscribers.

Parameters *msap-policy-name* — Specifies the MSAP policy for which all associated idle MSAPs are cleared.

msap-policy

Syntax **msap-policy** *msap-policy-name* [**idle-only**]

Context clear>subscr-mgmt

Description This command deletes Managed SAPs (MSAPs) created by the MSAP policy.

This command can remove an MSAP with active subscribers still associated with the MSAP. Use the **idle-only** parameter to remove only MSAPs in an idle state.

Parameters *msap-policy-name* — Specifies an existing managed SAP policy name. A string up to 32 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

idle-only — Specifies to remove only the MSAPs in an idle state.

peakvalue-stats

Syntax **peakvalue-stats iom** (*slot* | **all**) [**recursive**]
peakvalue-stats mda (*mda* | **all**) [**recursive**]
peakvalue-stats port (*port-id* | **all**)
peakvalue-stats pw-port (*pw-port* | **all**)
peakvalue-stats system [**recursive**]

Context clear> subscriber-mgmt

Description This command resets the most recent peak counter.

Note: Clearing one counter will not impact other counters. For example, clearing one IOM's most recent peak value will not impact chassis peak value.

Parameters *slot* — Clears IOM host peak value statistics for the specified IOM.

mda — Clears MDA host peak value statistics for the specified MDA.

port-id — Clears port host peak value statistics for the specified port ID.

pw-port — Clears pseudowire port host peak value statistics for the specified port.

Values 1 to 10239

system — Clears system host peak value statistics.

all — Clears all host peak value statistics.

recursive — Resets the sub-level counters. For example, clearing IOM counters with the **recursive** keyword also clears counters of all ports counters on that IOM.

radius-accounting

Syntax **radius-accounting** [*policy-name*]

Context clear> subscriber-mgmt

Description This command clears RADIUS accounting data for the specified policy.

Parameters *policy-name* — The name of the policy. The string is case sensitive and limited to 32 ASCII 7-bit printable characters with no spaces.

statistics

Syntax **statistics sub-profile** {*sub-profile-name* | **all**}
statistics sla-profile {*sla-profile-name* | **all**}

Context clear> subscriber-mgmt

Description This command clears the current and historical statistics.

Parameters *sub-profile-name* — Clears the statistics for a specific subscriber profile, limited to 32 characters.

sub-profile all — Clears all subscriber profile statistics.

sla-profile-name — Clears the statistics for a specific SLA profile, limited to 32 characters.

sla-profile all — Clears all SLA profile statistics.

scheduler-stats

Syntax **scheduler-stats**

Context clear>qos

Description This command clears scheduler statistics.

subscriber

Syntax **subscriber** *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* [**scheduler** *scheduler-name*] [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]
subscriber *sub-ident-string* [**scheduler** *scheduler-name*] [**ingress** | **egress**]

Context clear>qos>scheduler-stats

Description This command clears scheduler stats by subscriber.

Parameters *sub-ident-string* — Clears information for the specified subscriber name.

sap-id — Clears information for the specified SAP.

sla-profile-name — Clears information for the specified SLA profile name.

scheduler-name — Clears information for the specified scheduler.

spi-sharing-type — Clears information for sessions with the specified SPI sharing type.

Values per-group, per-session-ipoe, per-session-ppp

spi-sharing-id — Clears information for sessions with the specified SPI sharing type and sharing identifier. The SPI group identifier is:

- an identifier when the sharing type is per-group
- an internal PPP or IPoE sharing identifier when the sharing type is a per-PPP or per-IPoE session

Values 0 to 4294967295

egress — Clears egress information for the subscriber.

ingress — Clears ingress information for the subscriber.

srrp

Syntax **srrp**

Context clear>router

Description This command enables the context to clear and reset SRRP virtual router instances.

interface

Syntax **interface** *subscriber-interface* [**id** *srrp-id*]

Context clear>router>srrp

Description This command clears and resets SRRP interface instances.

Parameters *subscriber-interface* — Specifies an existing subscriber interface name up to 32 characters.

srrp-id — Specifies an existing SRRP ID.

Values 1 to 4294967295

statistics

Syntax **statistics interface** *subscriber-interface* [**id** *srrp-id*]

Context clear>router>srrp

Description This command clears statistics for SRRP instances.

Parameters *subscriber-interface* — Specifies an existing subscriber interface name up to 32 characters.

srrp-id — Specifies an existing SRRP ID.

Values 1 to 4294967295

route-downloader

Syntax **route-downloader** *name* [**vprn** *vprn*] [**family** *family*]

Context clear>aaa

Description This command clears all the radius-downloaded routes from the internal downloader cache (or protocol RIB/db) (and thus eventually from the RTM itself). The parameters **vprn** and/or **family** allow to restrict the deletion of those routes learned in a particular address family (IPv4 or IPv6) and/or a particular VPRN.

By default, all VPRNs and both IPv4 and IPv6 families are affected.

Note: A clear of the internal protocol DB means the corresponding prefix that were deleted should be removed from the RTM (and from any other exports) as well.

Parameters *vprn* — Specifies to limit the removal of prefixes to only the specific VPRN. The parameter can be either the service-id or service-name that identifies a VPRN.

family — Specifies to limit he removal or prefixes only belonging to the address family IPv4 or IPv6. Only these two values are accepted.

Values ipv4, ipv6

vport

Syntax **port** *port-id* **vport** *name* [**scheduler** *scheduler-name*]

Context clear>qos>scheduler-stats

Description This command clears the Vport scheduler stats.

Parameters *port-id* — Clears information for the specified port.

name — Clears information for the specified Vport.

scheduler-name — Clears information for the specified scheduler name.

ipoe session

Syntax **ipoe session** [**sap** *sap-id*] [**interface** *ip-int-name* | *ip-address*] [**mac** *ieee-address*] [**circuit-id** *circuit-id*] [**remote-id** *remote-id*] [**inter-dest-id** *intermediate-destination-id*] [**no-inter-dest-id**] [**ip-address** *ip-prefix*[/*prefix-length*]] [**port** *port-id*] [**subscriber** *sub-ident-string*] [**sap-session-id** *sap-session-index*]

ipoe session all

Context clear>service>id

Description This commands clears all identified IPoE sessions for the specified service instance. All associated subscriber hosts are deleted from the system.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

ip-int-name — Specifies the name of the IP interface up to 32 characters.

ieee-address — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

circuit-id — Specifies the circuit ID up to 127 characters.

remote-id — Specifies that information that goes into the remote-id sub-option in the DHCP relay packet, up to 255 characters.

ip-address[/prefix-length] — Specifies information for the specified IP address and mask.

no-inter-dest-id — Displays the information about no intermediate destination ID.

intermediate-destination-id — Specifies the intermediate destination identifier which is encoded in the identification strings up to 32 characters.

port-id — Displays information about the specified port ID.

sap-session-index — Displays sap-session-index information.

service-id — Specifies the service ID of the wholesaler. This parameter applies only to the 7750 SR.

Values

service-id:	1 to 2147483647
svc-name:	64 characters maximum

sub-ident-string — Specifies an existing subscriber identification profile.

service-id — Specifies information for the specified service ID.

Values

service-id:	1 to 2147483647
svc-name:	64 characters maximum

detail — Displays all IPoE session details.

all — clears all active IPoE sessions for the specified service instance.

msap

Syntax **msap** {*msap-id* | **idle-only**}

Context clear>service>id

Description This command clears Managed SAP (MSAP) information.

This command can remove an MSAP with active subscribers still associated with the MSAP. You can specify the **idle-only** parameter to clear only MSAPs in an idle state; MSAPs with active subscribers are not cleared.

Parameters *msap-id* — Specifies a Managed SAP ID.

Values

dot1q	[port-id | lag-id]:qtag1
qinq	[port-id | lag-id]:qtag1.qtag2
qtag1	0 to 4094
qtag2	0 to 4094

idle-only — Specifies to clear only MSAPs in an idle state.

msap-policy

Syntax **msap-policy** *msap-policy-name*

Context clear>service>id

Description This command clears Managed SAPs created by the Managed SAP policy.

Parameters *msap-policy-name* — Clears an existing MSAP policy.

statistics

Syntax statistics

Context clear>service

Description This command clears the statistics for a service.

subscriber

Syntax **subscriber** *sub-ident-string*

Context clear>service>statistics

Description This command clears the statistics for a particular subscriber.

Parameters *sub-ident-string* — Clears statistics for the specified subscriber identification string.

9.25.2.4 Tools Commands

tools

Syntax **tools**

Context root

Description The context to enable useful tools for debugging purposes.

Parameters **dump** — Enables dump tools for the various protocols.

perform — Enables tools to perform specific tasks.

avp-match-learned-session-id

Syntax **avp-match-learned-session-id**

Context tools>dump>debug>diameter

Description This command displays the diameter session IDs that are learned by the AVP value matching in a diameter debug.

Per **avp-match** *id* in a diameter peer policy, a single diameter session ID can be learned. "N/A" means no session ID is learned.

Output The following is a sample output for the **avp-match-learned-session-id** command.

Sample output

```
# tools dump debug diameter avp-match-learned-session-id
===============================================================================
Learned Session-Id
===============================================================================
Diameter Peer Policy             ID Session-Id
-------------------------------------------------------------------------------
diameter-peer-policy-1           1  bng.domain.com;1452197100;13
```

```
diameter-peer-policy-1           2  N/A
-------------------------------------------------------------------------------
No. of AVP matches: 2
===============================================================================
```

perform

Syntax **perform**

Context tools

Description This command enables the context to enable tools to perform specific tasks.

credit-reset

Syntax **credit-reset sap** *sap-id* **subscriber** *sub-ident-string* **sla-profile** *sla-profile-name* {**category** *category-name* | **all-categories**} [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]
credit-reset sap *sap-id* **ip** *ip-address* {**category** *category-name*| **all-categories**} [**mac** *mac-address*]
credit-reset svc *service-id* **ip** *ip-address* {**category** *category-name* | **all-categories**} [**mac** *mac-address*]

Context tools>perform>subscr-mgmt

Description This command resets the credit for an SLA-profile instance.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

sub-ident-string — Resets the credit for the specified subscriber identification, up to 32 characters.

sla-profile-name — Resets the credit for the specified SLA profile, up to 32 characters.

category-name — Resets the credit for the specified category, up to 32 characters.

spi-sharing-type — Resets the credit for sessions with the specified SPI sharing type.

Values per-group, per-session-ipoe, per-session-ppp

spi-sharing-id — Resets the credit for sessions with the specified SPI sharing type and sharing identifier. The SPI sharing identifier is:

- an identifier when the sharing type is per-group
- an internal PPP or IPoE identifier when the sharing type is a per-PPP or per-IPoE session

Values 0 to 4294967295

ip-address — Resets the credit for the specified IP address.

mac-address — Resets the credit for the specified MAC IP address.

sub-profile-string — Resets the credit for the specified subscriber profile, up to 32 characters.

service-id — Resets the credit for the specified service ID.

Values service-id:1 2147483647
svc-name: up to 64 characters.

persistence

Syntax **persistence**

Context tools>perform

Description This command enables the context to configure downgrade parameters.

downgrade

Syntax **downgrade target-version** *target* [**reboot**]

Context tools>perform>persistence

Description This command downgrades persistence files to a previous version.

Parameters *target* — Specifies the downgrade version.

reboot — Specifies to reboot the system after a successful conversion.

eval-msap

Syntax **eval-msap** {**policy** *msap-policy-name* | **msap** *sap-id*}

Context tools>perform>subscr-mgmt

Description This command re-applies the managed SAP policy to the managed SAP identified by the specified *sap-id* or to all managed SAPs associated with the specified *msap-policy name*.

Parameters *msap-policy-name* — Specifies an existing MSAP policy. The MSAP policy is re-applied to all associated managed SAPs.

sap-id — Specifies an MSAP SAP ID. The MSAP policy is re-applied to the specified MSAP.

Values [*port-id* | lag-*id*]:*qtag1*
[*port-id* | lag-*id*]:*qtag1.qtag2*

eval-lease-state

Syntax **eval-lease-state** [**svc-id** *service-id*] [**sap** *sap-id*] [**subscriber** *sub-ident-string*] [**ip** *ip-address*] [**mac** *mac-address*]

Context tools>perform>subscr-mgmt

Description This command evaluates lease state.

Parameters **sap-id** — Specifies the physical port identifier portion of the SAP definition.

ip-address — Specifies the a server's IP address. This address must be unique within the subnet and specified in dotted decimal notation. Allowed values are IP addresses in the range 1.0.0.0 – 223.255.255.255 (with support of /31 subnets).

sub-ident-string — Specifies the subscriber ID string, up to 32 characters.

service-id — Specifies an existing service ID.

Values 1 to 2147483647

subscriber-mgmt

Syntax **subscriber-mgmt**

Context tools>perform

Description This command enables the context to configure tools to control subscriber management.

eval-rates

Syntax **eval-rates**

Context tools>perform>subscr-mgmt>bonding

Description This command evaluates all bonding contexts and recalculates all reference rates. This command can be useful in handling changes that are not automatically tracked, such as adapting a scheduler rate in a QoS policy.

edit-lease-state

Syntax **edit-lease-state sap** *sap-id* **ip** *ip-address* [**mac** *ieee-address*] [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*] [**app-profile-string** [*app-profile-string*] [**inter-dest-id** *intermediate-destination-id*] [**ancp-string** *ancp-string*]

edit-lease-state svc-id *service-id* **ip** *ip-address* [**mac** *ieee-address*] [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*] [**app-profile-string** [*app-profile-string*] [**inter-dest-id** *intermediate-destination-id*] [**ancp-string**

ancp-string]

Context tools>perform>subscr-mgmt

Description This command provides the parameters to edit lease state information.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

ip-address — Modifies lease state information for the specified IP address.

ieee-address — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

sub-ident-string — Modifies the lease state information for the specified subscriber identification.

sub-profile-string — Modifies lease state information for the specified subscriber profile.

sla-profile-string — Modifies lease state information for the specified SLA profile.

service-id — Modifies lease state information for the specified service ID.

Values

service-id:	1 to 2147483647
svc-name:	64 characters maximum

coa

Syntax **coa** [**nas-port-id** *attr-string*] [**framed-ip-addr** *ip-address*] [**alc-ipv6-addr** *ipv6-address*] [**delegated-ipv6-pfx** *ipv6-prefix*[/*ipv6-prefix-length*]] [**framed-ipv6-pfx** *ipv6-prefix* [/*ipv6-prefix-length*]] [**alc-serv-id** *serv-id*] [**acct-session-id** *attr-string*] [**alc-subscr-id** *attr-string*] [**alc-brg-id** *attr-string*] [**alc-client-hw-addr** *mac-address*] [**attr** *attribute* [*attribute*]] [**attr-from-file** *file-url*] [**from-server** *attr-string*] [**router-or-service** *attr-string*] [**debug**]

Context tools>perform>subscr-mgmt

Description This command triggers a Change of Authorization (CoA) without a RADIUS authentication policy.

This command can be used to spoof a CoA from a configured server for purposes such as testing CoA python scripts. However, spoofing a CoA from a RADIUS server requires the configuration of a RADIUS authentication policy.

Parameters *attr-string* — Specifies the physical access circuit of the NAS up to 253 characters.

ip-address — Specifies the IPv4 host address.

Values

ipv4-address:	a,b,c,d

ipv6-address — Specifies the IPv6 host address.

Values

ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D

ipv6-prefix[/*ipv6-prefix-length*] — Specifies the IPv6 PD host prefix.

Values

ipv6-prefix:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
ipv6-prefix-length	[0 to 128]

ipv6-prefix[/*ipv6-prefix-length*] — Specifies the IPv6 SLAAC host prefix.

Values

ipv6-prefix:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
ipv6-prefix-length	[0 to 128]

serv-id — Specifies the service ID of the LNS subscriber.

Values 0 to 4294967295

attr-string — Specifies the subscriber accounting session ID up to 22 characters.

attr-string — Specifies the subscriber ID up to 32 characters.

attr-string — Specifies the BRG ID to a maximum of 32 characters.

mac-address — Specifies the subscriber host MAC address in the form: xx:xx:xx:xx:xx:xx.

attribute — Specifies the attributes (RADIUS VSAs) for override. The VSA must be specified in number format. For example, VSA Alc-Subsc-Prof-Str=”1M-prof” is expressed as 6527.12 =“1M-prof”. Up to five attributes can be specified within this command.

For more information on VSAs, refer to the *7750 SR and VSR RADIUS Attributes Reference Guide*.

Values

<type>=<value>

<type>:<tag>=<value>

<vendor>,<type>=<value>

e, <type>,<exttyp>=<value> (type=[241|242|243|244])

le,<type>,<exttyp>=<value> (type=[245|246])

evs,<type>,<vendor>,<vendor-type>=<value>

levs,<type>,<vendor>,<vendor-type>=<value>

file-url — Specifies the file URL where the VSAs are located.

Values

local-url	[*cflash-id/*] [*file-path*] (200 characters maximum including cflash-id directory length of 99 characters maximum each)
remote-url	[{**ftp**://|**tftp**://}*login*:*pswd*@*remote-locn*/][*file-path*] (255 characters maximum with directory length 99 characters maximum each)
remote-locn	[*hostname*	*ipv4-address*	*ipv6-address*]
ipv4-address:	a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x [-*interface*]
	x:x:x:x:x:x:d.d.d.d[-*interface*]
	x - [0 to FFFF]H
	d - [0 to 255]D
	interface - Specifies the link local addresses, up to 32 characters maximum.
cflash-id	cf1:	cf1-A:	cf1-B:	cf2 :| cf2-A:	cf2-B:	cf3:	cf3-A:	cf3-B:

attr-string — Specifies the RADIUS server name that is configured under the router or VPRN service context, to a maximum of 32 characters.

attr-string — Specifies the router or VPRN service instance where the RADIUS server is configured, to a maximum of 32 characters.

debug — Displays the debug message associated with the CoA command.

Output The following is an example of subscriber management CoA tools information.

Sample Output

```
Node# tools perform subscriber-mgmt coa alc-subscr-id subscriber-
01 attr 6527,13=new-sla-prof debug
===============================================================================
Inject Change-of-Authorization
===============================================================================
SUCCESS: got ACK
  Change of Authorization(43) id 0 len 50 from 0.0.0.0:3799 vrid 1
    VSA [26] 8 Alcatel(6527)
      SUBSC ID STR [11] 6 subscriber-01
    VSA [26] 10 Alcatel(6527)
      SLA PROF STR [13] 8 new-sla-prof
  Hex Packet Dump:
  1a 0e 00 00 19 7f 0b 08 77 69 66 69 2d 31 1a 10 00 00 19 7f 0d 0a 73 6c 61
  2d 70 72 6f 66
  Change of Authorization Ack(44) 0.0.0.0:3799 id 0 len 20 vrid 1
  Hex Packet Dump:
```

credit-reset

Syntax **credit-reset sap** *sap-id* **subscriber** *sub-ident-string* **sla-profile** *sla-profile-name* {**category** *category-name* | **all-categories**}
credit-reset sap *sap-id* **ip** *ip-address* {**category** *category-name* | **all-categories**}
credit-reset svc *service-id* **ip** *ip-address* {**category** *category-name* | **all-categories**}

Context tools>perform>subscr-mgmt

Description This command resets the credit for an SLA-profile instance.

Parameters *sap-id* — Specifies the physical port identifier portion of the SAP definition.

ip-address — Specifies lease state information for the specified IP address.

sub-ident-string — Specifies lease state information for the specified subscriber identification.

sub-profile-string — Specifies lease state information for the specified subscriber profile.

sla-profile-string — Specifies lease state information for the specified SLA profile.

service-id — Specifies lease state information for the specified service ID.

Values

service-id:	1 to 2147483647
svc-name:	64 characters maximum

edit-ppp-session

Syntax **edit-ppp-session sap** *sap-id* [**user-name** *user-name*] [**mac** *mac-address*] [**session-id** *session-id*] [**ip-address** *ip-address*] [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*] [**inter-dest-id** *intermediate-destination-id*] [**ancp-string** *ancp-string*] [**app-profile-string** *app-profile-string*]
edit-ppp-session svc-id *service-id* [**user-name** *user-name*] [**mac** *mac-address*] [**session-id** *session-id*] [**ip-address** *ip-address*] [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*] [**inter-dest-id** *intermediate-destination-id*] [**ancp-string** *ancp-string*] [**app-profile-string** *app-profile-string*]

Context tools>perform>subscr-mgmt

Description This command updates the data of the PPP session identified with the given MAC address and SAP identifier. Optionally the remote-id and circuit-id can be specified to identify the IPoE session to update.

Note: The changes take immediate effect.

Parameters *sap-id* — Specifies a SAP ID, up to 255 characters.

user-name — Specifies the PPP user name, up to 253 characters.

mac-address — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers.

session-id — Specifies information about the ID of the PPP session.

Values 1 to 65535

sub-ident-string — Specifies an existing subscriber identification profile.

sla-profile-string — Identifies the SLA profile string up, to 16 characters.

ip-address — Specifies the peer's IP address.

service-id — Specifies an existing subscriber service ID.

Values 1 to 2147483647

sub-profile-string — Specifies the subscriber profile string, up to 16 characters.

sub-profile-name — Specifies the subscriber profile name, up to 32 characters.

intermediate-destination-id — Specifies the intermediate destination identifier which is encoded in the identification strings, up to 32 characters.

ancp-string — Specifies the ASCII representation of the string, up to 63 characters.

app-profile-string — Specifies the application profile string, up to 16 characters.

edit-slaac-host

Syntax **edit-slaac-host sap** *sap-id* [**user-name** *user-name*] [**mac** *mac-address*] [**session-id** *session-id*] [**ip-address** *ip-address*] [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*] [**inter-dest-id** *intermediate-destination-id*] [**ancp-string** *ancp-string*] [**app-profile-string** *app-profile-string*]

edit-slaac-host svc-id *service-id* [**user-name** *user-name*] [**mac** *mac-address*] [**session-id** *session-id*] [**ip-address** *ip-address*] [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** *sla-profile-string*] [**inter-dest-id** *intermediate-destination-id*] [**ancp-string** *ancp-string*] [**app-profile-string** *app-profile-string*]

Context tools>perform>subscr-mgmt

Description This command updates the data of the SLAAC host information.

Parameters *sap-id* — Specifies a SAP ID, up to 255 characters maximum.

user-name — Specifies the PPP user name up to 253 characters.

mac-address — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers.

session-id — Specifies information about the ID of the PPP session.

Values 1 to 65535

sub-ident-string — Specifies an existing subscriber identification profile.

sla-profile-string — Identifies the SLA profile string up to 16 characters.

ip-address — Specifies the peer's IP address.

service-id — Specifies an existing subscriber service ID.

Values 1 to 2147483647

sub-profile-string — Specifies the subscriber profile string up to 16 characters.

sub-profile-name — Specifies the subscriber profile name up to 32 characters.

intermediate-destination-id — Specifies the intermediate destination identifier which is encoded in the identification strings up to 32 characters.

ancp-string — Specifies the ASCII representation of the up to 63 characters.

app-profile-string — Specifies the application profile string up to 16 characters.

edit-ipoe-session

Syntax **edit-ipoe-session sap** *sap-id* **mac** *mac-address* [**subscriber** *sub-ident-string*] [**sub-profile-string** *sub-profile-string*] [**sla-profile-string** sla-*profile-string*] [**inter-dest-id** *intermediate-destination-id*] [**ancp-string** *ancp-string*] [**app-profile-string** *app-profile-string*] [**circuit-id** *circuit-id*] [**remote-id** *remote-id*]

Context tools>perform>subscr-mgmt

Description This command updates the data of the IPoE session identified with the given MAC address and SAP identifier. Optionally the remote ID and circuit ID can be specified to identify the IPoE session to update.

Note: The changes take immediate effect.

eval-ipoe-session

Syntax **eval-ipoe-session** [**svc-id** *service-id*] [**sap** *sap-id*] [**mac** *mac-address*] [**circuit-id** *circuit-id*] [**remote-id** *remote-id*] [**subscriber** *sub-ident-string*]

Context tools>perform>subscr-mgmt

Description This command re-evaluates the mapping between authentication strings such as the SLA profile string and the actual profiles for the identified IPoE sessions.

Parameters *service-id* — Specifies an existing subscriber service ID.

Values 1 to 2147483647

sap-id — Specifies a SAP ID, up to 255 characters.

mac-address — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers.

circuit-id — Specifies the circuit ID up to 127 characters.

remote-id — Specifies that information that goes into the remote-id sub-option in the DHCP relay packet, up to 255 characters.

sub-ident-string — Specifies an existing subscriber identification profile.

eval-ppp-session

Syntax **eval-ppp-session** [**svc-id** *service-id*] [**user-name** *user-name*] [**sap** *sap-id*] [**ip** *ip-address*] [**subscriber** *sub-ident-string*]
eval-ppp-session [**svc-id** *service-id*] [**user-name** *user-name*] [**sap** *sap-id*] **mac** *mac-address* [**session-id** *session-id*] [**ip** *ip-address*] [**subscriber** *sub-ident-string*]

Context tools>perform>subscr-mgmt

Description This command evaluates PPP sessions.

Parameters *service-id* — Specifies an existing subscriber service ID.

Values 1 to 2147483647

user-name — Specifies the PPP user name up to 253 characters.

sap-id — Specifies a SAP ID up to 255 characters.

ip-address — Specifies the peer's IP address.

sub-ident-string — Specifies an existing subscriber identification profile.

mac-address — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers.

session-id — Specifies information about the ID of the PPP session.

Values 1 to 65535

eval-slaac-host

Syntax **eval-slaac-host** [**svc-id** *service-id*] [**sap** *sap-id*] [**subscriber** *sub-ident-string*] [**ipv6-address** *ipv6-prefix*] [**mac** *ieee-address*]

Context tools>perform>subscr-mgmt

Description This command evaluates the SLAAC host.

Parameters *service-id* — Specifies an existing subscriber service ID.

Values 1 to 2147483647

sap-id — Specifies a SAP ID, up to 255 characters maximum.

sub-ident-string — Specifies an existing subscriber identification profile.

ieee-address — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers.

ipv6-prefix — Specifies an IPv6 prefix and prefix length.

Values

ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d

remap-lease-state

Syntax **remap-lease-state old-mac** *ieee-address* **mac** *ieee-address*
remap-lease-state sap *sap-id* [**mac** *ieee-address*]

Context tools>perform>subscr-mgmt

Description This command allows the remapping of all existing hosts if network card on CMTS/WAC side is changed is required.

When this command is executed, the following restrictions apply:

- When **sap** is taken, all leases associated with the SAP are re-written.
  - For a SAP with a configured MAC in **lease-populate** command, this MAC is taken.
  - For a SAP without a configured MAC the MAC from tools command is taken.
  - For a SAP without a configured MAC and no MAC in tools command no action is performed.
- When using the **old-mac** option, providing a new MAC *ieee-address* is mandatory.

This command is applicable only when dealing with DHCP lease states which were instantiated using l2header mode of DHCP operation.

Parameters **old-mac** *ieee-address* — Specifies the old MAC address to remap.

mac *ieee-address* — Specifies that the provisioned MAC address is used in the anti-spoofing entries for this SAP when l2-header is enabled. The parameter may be changed mid-session. Existing sessions will not be re-programmed unless a **tools perform** command is issued for the lease.

sap-id — Specifies the physical port identifier portion of the SAP definition.

When configured, the SAP parameter remaps all MAC addresses of DHCP lease states on the specified SAP. When no optional MAC parameter is specified, the **sap** *sap-id* command remaps all MAC addresses of lease states towards the MAC address specified in the l2-header configuration.

re-ident-sub

Syntax **re-ident-sub** *old-sub-ident-string* **to** *new-sub-ident-string*

Context tools>perform>subscr-mgmt

Description This command renames a subscriber identification string.

Parameters *old-sub-ident-string* — Specifies the existing subscriber identification string to be renamed. With implicit subscriber ID generation enabled, the *old-sub-ident-string* cannot refer to a subscriber ID in an implicitly generated default format.

new-sub-ident-string — Specifies the new subscriber identification string name. With implicit subscriber ID generation enabled, the *new-sub-ident-string* cannot refer to a subscriber ID in an implicitly generated default format.

redundancy

Syntax **redundancy**

Context tools>dump

Description This command enables the context to dump redundancy parameters.

multi-chassis

Syntax **multi-chassis**

Context tools>dump>redundancy

Description This command enables the context to dump multi-chassis parameters.

mc-endpoint

Syntax **mc-endpoint peer** *ip-address*

Context tools>dump>redundancy>multi-chassis

Description This command dumps multi-chassis endpoint information.

Parameters *ip-address* — Specifies the peer's IP address.

force-switchover

Syntax **force-switchover tunnel-group** *local-group-id*

Context tools>dump>redundancy>multi-chassis

Description This command manually switches over mc-ipsec mastership of the specified tunnel-group.

Parameters *local-group-id* — Specifies the local tunnel-group ID configured under the **config>redundancy.multichassis>peer>mc-ipsec** context.

mc-ring

Syntax **mc-ring**
mc-ring peer *ip-address* [**ring** *sync-tag*]

Context tools>dump>redundancy>multi-chassis

Description This command dumps multi-chassis ring information.

Parameters *ip-address* — Specifies the peer's IP address.

sync-tag — Specifies the ring's sync-tag created in the **config>redundancy>mc>peer>mcr>ring** context.

srrp-sync-database

Syntax **srrp-sync-database** [**instance** *instance-id*] [**peer** *ip-address*]

Context tools>dump>redundancy>multi-chassis

Description This command dumps SRRP database information and only applies to the 7450 ESS and 7750 SR.

Parameters *ip-address* — Specifies the peer's IP address

instance *instance-id* — Dumps information for the specified Subscriber Router Redundancy Protocol instance configured on this system.

Values 1 to 4294967295

sync-database

Syntax **sync-database** [**peer** *ip-address*] [**port** *port-id* | *lag-id*] [**sync-tag** *sync-tag*] [**application** *application*] [**detail**] [**type** *type*]

Context tools>dump>redundancy>multi-chassis

Description This command dumps MCS database information.

Parameters *ip-address* — Specifies the peer's IP address.

port-id | *lag-id* — Indicates the port or LAG ID to be synchronized with the multi-chassis peer.

Values *slot/mda/port* or lag-*lag-id*

sync-tag — Specifies a synchronization tag to be used while synchronizing this port with the multi-chassis peer.

application — Specifies a particular multi-chassis peer synchronization protocol application.

Values

dhcp-server:	local dhcp server (7450 ESS and 7750 SR only)
igmp:	Internet group management protocol
igmp-snooping:	igmp-snooping
mc-ring:	multi-chassis ring
mld-snooping:	multicast listener discovery-snooping
srrp:	simple router redundancy protocol (7450 ESS and 7750 SR only)
sub-host-trk:	subscriber host tracking (7450 ESS and 7750 SR only)
sub-mgmt:	subscriber management (7450 ESS and 7750 SR only)

type — Indicates the locally deleted or alarmed deleted entries in the MCS database per multi-chassis peer.

Values alarm-deleted, local-deleted

detail — Displays detailed information.

mc-ipsec

Syntax **mc-ipsec**

Context tools>perform>redundancy>multi-chassis

Description This command enables the context to configure mc-ipsec parameters.

mc-endpoint

Syntax **mc-endpoint peer** *ip-address*

Context tools>dump>redundancy>multi-chassis

Description This command dumps multi-chassis endpoint information.

Parameters *ip-address* — Specifies the peer's IP address.

mc-ring

Syntax **mc-ring**
mc-ring peer *ip-address* [**ring** *sync-tag*]

Context tools>dump>redundancy>multi-chassis

Description This command dumps multi-chassis ring information.

Parameters *ip-address* — Specifies the peer's IP address.

sync-tag — Specifies the ring's sync-tag created in the **config>redundancy>mc>peer>mcr>ring** context.

srrp-sync-database

Syntax **srrp-sync-database** [**instance** *instance-id*] [**peer** *ip-address*]

Context tools>dump>redundancy>multi-chassis

Description This command dumps SRRP database information and only applies to the 7450 ESS and 7750 SR.

Parameters *ip-address* — Specifies the peer's IP address

instance-id — Dumps information for the specified Subscriber Router Redundancy Protocol instance configured on this system.

Values 1 to 4294967295

sync-database

Syntax **sync-database** [**peer** *ip-address*] [**port** *port-id* | *lag-id*] [**sync-tag** *sync-tag*] [**application** *application*] [**detail**] [**type** *type*]

Context tools>dump>redundancy>multi-chassis

Description This command dumps MCS database information.

Parameters *ip-address* — Specifies the peer's IP address.

port-id **|** *lag-id* — Indicates the port or LAG ID to be synchronized with the multi-chassis peer.

Values *slot/mda/port* or lag-*lag-id*

sync-tag — Specifies a synchronization tag to be used while synchronizing this port with the multi-chassis peer.

application — Specifies a particular multi-chassis peer synchronization protocol application.

Values

dhcp-server:	local dhcp server (7450 ESS and 7750 SR only)
igmp:	Internet group management protocol
igmp-snooping:	igmp-snooping
mc-ring:	multi-chassis ring
mld-snooping:	multicast listener discovery-snooping
srrp:	simple router redundancy protocol (7450 ESS and 7750 SR only)
sub-host-trk:	subscriber host tracking (7450 ESS and 7750 SR only)
sub-mgmt:	subscriber management (7450 ESS and 7750 SR only)

type — Indicates the locally deleted or alarmed deleted entries in the MCS database per multi-chassis peer.

Values alarm-deleted, local-deleted

detail — Displays detailed information.

forcerenew

Syntax **forcerenew svc-id** *service-id* {**ip** *ip-address*[*Imask*] | **mac** *ieee-address*}
forcerenew {**interface** *interface-name* | **sap** *sap-id* | **sdp** *sdp-id:vc-id*} [**ip** *ip-address*[*Imask*] |**mac** *ieee-address*]

Context tools>perform>subscr-mgmt

Description This command forces the renewal of lease state and only applies to the 7450 ESS and 7750 SR.

In cases where the IPoE hosts resides in a private retail subnet VPRN, the MAC address must be used for the **tools**>**perform**>**subscriber-mgmt**>**forcerenew** command instead of the IP address.

Parameters *service-id* — Forces renewal of the lease state for the specified service ID.

Values 1 to 2147483647

sap-id — Forces renewal of the lease state for the specified SAP.

ip-address — Forces renewal of the lease state for the specified IP address.

ieee-address — Forces renewal of the lease state for the specified MAC address.

interface-name — Forces renewal of the lease state for the specified interface name.

local-user-db

Syntax **local-user-db** *local-user-db-name*

Context tools>perform>subscr-mgmt

Description This command provides tools to control the local user database.

Parameters *local-user-db-name* — Specifies the name of a local user database up to 32 characters.

ipoe

Syntax **ipoe**

Context tools>perform>subscr-mgmt>loc-user-db

Description This command provides the tools to control IPoE entries in the local user database.

host-lookup

Syntax **host-lookup** [**mac** *ieee-address*] [**remote-id** *remote-id*-ascii] [**sap-id** *sap-id*] [**service-id** *service-id*] [**string** *vso-string*] [**system-id** *system-id*] [**option60** *option-60-ascii*] [**circuit-id** *circuit-id-ascii*] [**circuit-id-hex** *circuit-id-hex*] [**option60-hex** *option60-hex*] [**remote-id-hex** *remote-id-hex*] [**derived-id** *derived-id*] [**ip-prefix** *ip-prefix/ip-prefix-length*]

Context tools>perform>subscr-mgmt>loc-user-db>ipoe

Description This command performs a lookup in the local user database.

Parameters *ieee-address* — Specifies the 48-bit MAC address for the static ARP in the form aa:bb:cc:dd:ee:ff or aa-bb-cc-dd-ee-ff where aa, bb, cc, dd, ee, and ff are hexadecimal numbers. Allowed values are any non-broadcast, non-multicast MAC and non-IEEE reserved MAC addresses.

remote-id — Specifies the information that goes into the **remote-id sub-option** in the DHCP relay packet up to 255 characters.

remote-id-ascii — Specifies that information that goes into the **remote-id sub-option** in the DHCP relay packet, up to 255 ACII characters.

sap-id — Specifies a SAP ID, up to 255 characters.

service-id — Specifies an existing subscriber service ID.

Values 1 to 2147483647

vso-string — Specifies a Vendor Specific Option (VSO) string.

system-id — Specifies the system ID up to 255 characters.

option-60-ascii — Specifies the Vendor-Identifying Vendor Option to match against. Option 60 is encoded as Type - Length - Value and must be formatted in ASCII.

option-60-hex — Specifies the Vendor-Identifying Vendor Option to match against. Option 60 is encoded as Type - Length - Value and must be formatted in HEX

ppp

Syntax **ppp**

Context tools>perform>subscr-mgmt>loc-user-db

Description This command provides the tools to control PPPoE entries in the local user database.

authentication

Syntax **authentication user-name** *user-name* **password** *password* [**circuit-id** *circuit-id*] [**circuit-id-hex** *circuit-id-hex*] [**derived-id** *derived-id*] [**mac** *ieee-address*] [**remote-id** *remote-id*] [**remote-id-hex** *remote-id-hex*] [**sap-id** *sap-id*] [**service-name** service-name]

Context tools>perform>subscr-mgmt>loc-user-db>ppp

Description This command authenticates the PPPoE user name.

Parameters *user-name* — Specifies the PPP user name up to 253 characters.

password — Specifies the password of this host up to 64 characters.

circuit-id — Specifies the circuit ID up to 127 characters.

circuit-id-hex — Specifies the circuit ID in hexadecimal format.

Values 0x0 to 0xFFFFFFFF (maximum 254 hex nibbles)

derived-id — Specifies an ASCII string that uniquely identifies a host, and is derived by a Python script from packet content available during a DHCP transaction up to 255 characters.

ieee-address — Specifies information about the MAC address of the PPP session.

remote-id — Specifies that information that goes into the remote-id sub-option in the DHCP relay packet, up to 255 characters.

remote-id-hex — Specifies the remote ID in hexadecimal format.

Values 0x0 to 0xFFFFFFFF (maximum 510hex nibbles)

sap-id — Specifies a SAP ID, up to 255 characters.

service-name — Specifies the service name.

host-lookup

Syntax **host-lookup** [**circuit-id** *circuit-id*] [**circuit-id-hex** *circuit-id-hex*] [**derived-id** *derived-id*] [**mac**

ieee-address] [**remote-id** *remote-id*] [**remote-id-hex** *remote-id-hex*] [**sap-id** *sap-id*] [**service-name** *service-name*] [**user-name** *user-name*]

Context tools>perform>subscr-mgmt>loc-user-db>ppp

Description This command performs a lookup in the local user database.

Parameters *circuit-id circuit-id* — Specifies the circuit ID from the Option 82.

circuit-id-hex — Specifies the circuit ID in hexadecimal format from the Option 82.

Values 0x0 to 0xFFFFFFFF (maximum 254 hex nibbles)

derived-id — Specifies an ASCII string that uniquely identifies a host, and is derived by a Python script from packet content available during a DHCP transaction.

remote-id — Specifies that information that goes into the remote-id sub-option in the DHCP relay packet, up to 255 characters.

remote-id-hex — Specifies the remote ID in hexadecimal format.

Values 0x0 to 0xFFFFFFFF (maximum 510hex nibbles)

sap-id — Specifies a SAP ID, up to 255 characters.

service-name — Specifies a PPP service name up to 253 characters.

user-name — Specifies a user name up to 128 characters.

mc-ring

Syntax **mc-ring**

Context tools>dump>redundancy>multi-chassis

Description This command dumps multi-chassis ring data.

sync-database

Syntax **sync-database** [**peer** *ip-address*] [**port** *port-id* | *lag-id*] [**sync-tag** *sync-tag*] [**application** *application*] [detail] [type *type*]

Context tools>dump>redundancy>multi-chassis

Description This command dumps multi-chassis sync database information.

Parameters *ip-address* — Dumps the specified address of the multi-chassis peer

port-id — Dumps the specified port ID of the multi-chassis peer

lag-id — Dumps the specified Link Aggregation Group (LAG) on this system

sync-tag — Dumps the synchronization tag used while synchronizing this port with the multi-chassis peer

application — Dumps the specified application information that was synchronized with the multi-chassis peer

Values dhcps, igmp, igmp-snooping, mc-ring, srrp, sub-mgmt, mld-snooping, all

detail — Displays detailed information

type — Filters by the specified entry type

Values alarm-deleted, local-deleted

sync-database-reconcile

Syntax **sync-database-reconcile** [**peer** *ip-address*] [**port** *port-id* | *lag-id* [**sync-tag** *sync-tag*]] [**application** *application*]
sync-database-reconcile [**peer** *ip-address*] [**sdp** *sdp-id [***sync-tag** *sync-tag]]* [**application** *application*]

Context tools>perform>redundancy>multi-chassis>mc-ipsec

Description This command provides the parameters to reconcile MCS database entries.

Parameters *ip-address* — Specifies the IP address (in the form of a.b.c.d).

port-id — Specifies the port ID in the slot/MDA/port format.

lag-id — Specifies the LAG ID.

Values lag-id: **lag**-*id*
lag: keyword
id: 1 to 200

syn-tag — Specifies the syn tag up to 32 characters.

application — Specifies the application.

Values dhcp-server — Specifies the local DHCP server
igmp — Specifies the Internet group management protocol
igmp-snooping — Specifies igmp-snooping
mc-ring — Specifies multi-chassis ring
l2tp — Specifies L2TP
mld — Specifies multicast listener discovery
mld-snooping — Specifies multicast listener discovery-snooping
srrp — Specifies simple router redundancy protocol
sub-host-trk — Specifies subscriber host tracking
sub-mgmt-ipoe — Specifies subscriber management for IPoE
sub-mgmt-pppoe — Specifies subscriber management for PPPoE
mc-ipsec — Specifies multi-chassis IPsec
python — Specifies Python cache

diameter-proxy — Specifies diameter proxy

pim-snpg-sap — Specifies protocol independent multicast snooping for SAPs

pim-snpg-sdp — Specifies protocol independent multicast snooping for SDPs

sdp-id — Specifies the SDP ID.

Values 1 to 17407

srrp-sync-data

Syntax **srrp-sync-database** [**instance** *instance-id*] [**peer** *ip-address*]

Context tools>dump>redundancy>multi-chassis

Description This command dumps multi-chassis SRRP sync database information.

Parameters *instance-id* — Specifies the instance ID.

Values 1 to 4294967295

ip-address — Dumps the specified address (in the form of a.b.c.d).

route-downloader

Syntax **route-downloader start** [**force**]

Context tools>perform>aaa

Description This command causes the download process to start immediately. If an ongoing download is already in progress then no further action is needed, except if the **force** keyword is added. In case the **force** keyword is added, then the current download is aborted and a new one is immediately restarted. If aborting the current download, the internal route table should not be emptied or cleared.

Parameters **start** — Starts the download process immediately.

force — Causes the current download to be aborted and a new one is immediately restarted.

router

Syntax **router** [*router-instance*]
router service-name *service-name*

Context tools>dump

Description This command enables the context for tools dump commands in the specified routing instance.

Parameters *router-instance* — Specifies the router instance name (Base or management) or service ID.

service-name — Specifies the configured service name for the routing instance.

ipoe-session

Syntax **ipoe-session**

Context tools>dump>router

Description This command enables the context for IPoE session-related tools dump commands in the specified routing instance.

migration

Syntax **migration** [**interface** *ip-int-name*]

Context tools>dump>router>ipoe-session

Description This command displays details on the IPoE session migration progress. It shows per group interface the number of hosts per type (DHCPv4, DHCPv6 and SLAAC) that are associated with an IPoE session or that are not associated with an IPoE session.

Output The following is an example of IPoE session migration information.

Sample Output

```
# tools dump router ipoe-session migration interface group-int-1-1
===============================================================================
 Type                Total     IPoE session     Non IPoE session
===============================================================================
 Group-interface: group-int-1-1 (IPoE session enabled)
-------------------------------------------------------------------------------
 DHCPv4              2         1                1
 DHCPv6              3         2                1
 SLAAC               0         0                0
-------------------------------------------------------------------------------
 IPoE sessions       1
===============================================================================
```

node

Syntax [**no**] **node** *host-name*

Context debug>diameter

Description This command debugs the Diameter node. Node-level debugging can report on all message exchange between the peers. Although this level can report messages that contain session id (app level messages), this level is session unaware. It deals strictly with getting the messages in and out of the system (connection level messages which are not routable, and application level messages which are routable).

Parameters *host-name* — Specifies the host name, up to 80 characters

on-error

Syntax [**no**] **on-error**

Context debug>diameter>node

Description This command debugs Diameter node errors. Only errors raised before the peer is determined. For example, there is no route for this message in the realm routing table, and therefore the peer cannot be determined.

peer

Syntax [**no**] **peer** *destination-host*

Context debug>diameter>node

Description This command debugs Diameter node peers. At this level, the forwarding/routing phase is completed and the peer is known. All messages flowing between this node and the peer are reported. Although the messages displayed can contain session-ids, this debugging level is session unaware (the session states are not maintained at this level).

Parameters *destination-host* — Specifies the host name, up to 80 characters

on-error

Syntax [**no**] **on-error**

Context debug>diameter>node>peer

Description This command reports only peer error conditions.

peer-to-peer

Syntax [**no**] **peer-to-peer**

Context debug>diameter>node>peer

Description This command reports only peer level message that are required for bringing up, maintaining and tearing down the peering connection (CER/A, DWR/A, and so on).

9.25.2.5 Debug Commands

diameter

Syntax [**no**] **diameter**

Context debug

Description This command enables debugging for diameter.

The **no** form of this command disables debugging.

application

Syntax **application**

Context debug>diameter

Description This command debugs application processing for the Diameter node. This level is session aware (the session state is maintained at this level). Connection level messages are not reported on this level.

on-error

Syntax [**no**] **on-error**

Context debug>diameter>application
debug>diameter>application>policy

Description This command debugs Diameter application errors and reports only peer error conditions, for example, an unknown session-id.

policy

Syntax [**no**] **policy** *policy-name*

Context debug>diameter>application

Description This command debugs Diameter applications for a particular application policy.

Parameters *policy-name* — Specifies the policy name, up to 32 characters

session-messages

Syntax [**no**] **session-messages**

Context debug>diameter>application>policy

Description This command debugs session messages.

dest-realm

Syntax **dest-realm** *realm*
no dest-realm

Context debug>diameter

Description This command restricts the output to a specific destination-realm.

The **no** form of this command disables debugging.

Parameters *realm* — Specifies the realm up to 80 characters.

detail-level

Syntax **detail-level** *level*

Context debug>diameter

Description This command configures the detail level of debug output.

The **no** form of this command disables debugging.

Parameters *level* — Specifies the detail level.

Values low, medium, high

diameter-peer

Syntax **diameter-peer peer** [*psm-events*]
no diameter-peer

Context debug>diameter

Description This command restricts output to a specific peer.

The **no** form of this command disables debugging.

Parameters *psm-events* — Specifies to restrict output to the peer's state machine (PSM).

diameter-peer-policy

Syntax **diameter-peer-policy** *policy*
no diameter-peer-policy *policy*

Context debug>diameter

Description This command restricts the debug output to messages on peers of a specific diameter peer policy. Up to eight diameter peer policies can be specified.

The **no** form of this command removes the specified diameter peer policy and all debug statement under this policy from the debug configuration.

Parameters *policy* — Specifies the diameter peer policy name, up to 32 characters.

avp-match

Syntax **avp-match** *id*
no avp-match *id*

Context debug>diam>diam-peer-plcy

Description This command restricts the debug output to messages within the diameter peer policy that belong to a diameter session identified based on the AVP value matching in a diameter application message.

At least the message type and one AVP match criteria must be specified in an **avp-match** *id* command.

If a diameter application message matches all criteria within one AVP match ID, then the session ID is learned and all subsequent messages of that diameter session are shown until a relearning occurs. (OR function between **avp-match** *id* commands.)

When the session ID is learned in an Answer message, an attempt is made to include the corresponding Request message in the debug output: The Request message should still be available in the system and must pass all debug filters (such as **message-type**).

By default an **avp-match** *id* is disabled and must be configured with the **debug>diam>diam-peer-plcy>avp-match no shutdown** to activate.

Parameters *id* — Specifies the AVP match ID, up to five per diameter peer policy.

Values 1 to 5

avp

Syntax **avp** *avp-id* **type** *type* [**ascii** | **hex**] **value** *value*
no avp *avp-id*

Context debug>diam>diam-peer-plcy>avp-match

Description This command specifies an AVP match criteria for AVP value matching. At least one and up to five AVP match criteria can be specified in an **avp-match** *id* command. When multiple AVP match criteria are specified, they must all match to be successful and result in a diameter session ID learning. (AND function between **avp** *avp-id* commands.)

The AVP in an AVP match criteria is identified by its AVP ID. The AVP ID is specified as [*vendor-id*-]*avp-code*[.*avp-id*] with nesting up to five levels deep.

The format type of the AVP should match the standard documents in which the AVP is specified. Any AVP can be specified as an octet string in hex format.

Parameters *avp-id* — Specifies that an AVP is a [*vendor-id*-]*avp-code*[.*avp-id*], up to five levels deep. For example to specify the Multiple-Services-Credit-Control.Quota-Holding-Time AVP, use *avp-id* = 456.10415-871.

Values *vendor-id*: [1 to 4294967295]
avp-code: [1 to 4294967295]

type — Specifies the format type of the AVP as octets string, integer32, integer64, unsigned32, unsigned64 or address. The type should match the format as specified in the corresponding standard documents. Any AVP can be specified as octet string in hex.

ascii | hex — Specifies the value format. The ASCII format is default. The hexadecimal format value should start with "0x". For example: 0x0000000a to specify a 4 byte integer value in hexadecimal format.

value — Specifies the actual value of the AVP that should be matched against (256 characters maximum; 127 characters maximum when specified in hex).

message-type

Syntax **message-type** [**ccr**] [**cca**] [**rar**] [**raa**] [**asr**] [**asa**] [**aar**] [**aaa**]
message-type all
no message-type

Context debug>diam>diam-peer-plcy>avp-match

Description This command specifies the message type match criteria for AVP value matching. Only specified diameter application messages are used for AVP value matching. This is a mandatory criteria in an **avp-match** *id* command.

This command does not restrict the debug output to the specified messages.

Parameters **ccr** — credit control request

cca — credit control answer

rar — re-authentication request

raa — re-authentication answer

asr — abort session request

asa — abort session answer

aar — aa request

aaa — aa answer

all — all message types

shutdown

Syntax [**no**] **shutdown**

Context debug>diam>diam-peer-plcy>avp-match

Description This command enables or disables the **avp-match** *id* criteria for filtering debug output based on AVP value matching.

A shutdown of the **avp-match** *id* clears the learned diameter session ID.

diameter-peer

Syntax **diameter-peer** *peer* [**psm-events**]
no diameter-peer

Context debug>diam>diam-peer-plcy

Description This command restricts the debug output to messages on a specific peer within the diameter peer policy. Overrides the **debug>diam>diam-peer** debug command for the specified diameter peer policy. A single peer can be specified per diameter peer policy.

When no peer is specified at the diameter peer policy level, the debug output is restricted to the peer configured at the **debug>diam** level.

The **no** form of this command removes the peer from the **debug>diam>diam-peer-plcy** debug output.

Parameters *peer* — Specifies a peer name, up to 32 characters.

psm-events — Specifies to include Peer State Machine (PSM) events in the debug output.

message-type

Syntax **message-type** [**ccr**] [**cca**] [**cer**] [**cea**] [**dwr**] [**dwa**] [**dpr**] [**dpa**] [**rar**] [**raa**] [**asr**] [**asa**] [**aar**] [**aaa**]
message-type all
no message-type

Context debug>diam>diam-peer-plcy
debug>diam

Description This command restricts the debug output to the specified message types.

When specified within a diameter peer policy, it overrides the message type configuration at the **debug>diam** level for messages received and sent on that diameter peer policy.

The **no** form of this command removes the message type from the debug configuration.

Parameters **ccr** — credit control request

cca — credit control answer

cer — capabilities exchange request

cea — capabilities exchange answer

dwr — device watchdog request

dwa — device watchdog answer

dpr — disconnect peer request

dpa — disconnect peer answer

rar — re-authentication request

raa — re-authentication answer

asr — abort session request

asa — abort session answer

aar — aa request

aaa — aa answer

all — all message types

origin-realm

Syntax **origin-realm** *realm*
no origin-realm

Context debug>diam

Description This command restricts output to a specific origin-realm.

arp-host

Syntax [**no**] **arp-host**

Context debug>service>id

Description This command enables and configures ARP host debugging.

The **no** form of this command disables ARP host debugging.

one-time-http-redirection

Syntax **one-time-http-redirection**

Context debug>service>id

Description This command produces one-time HTTP redirection debug output.

ppp

Syntax [**no**] **ppp**

Context debug>service>id>

Description This command enables the PPP debug context.

The **no** form of this command disables PPP debugging.

event

Syntax [**no**] **event**

Context debug>service>id>ppp

Description This command enables the PPP event debug context.

The **no** form of this command disables PPP event debugging.

dhcp-client

Syntax **dhcp-client** [**terminate-only**]
no dhcp-client

Context debug>service>id>ppp>event

Description This command enable PPP event debug for DHCP client.

The **no** form of this command disables PPP event debugging for the DHCP client.

Parameters **terminate-only** — Enables debug for local terminated PPP session.

l2tp

Syntax **l2tp** [**terminate-only**]
no l2tp

Context debug>service>id>ppp>event

Description This command enables PPP L2TP event debug.

Parameters **terminate-only** — Enables debug for local terminated PPP session.

local-address-assignment

Syntax **local-address-assignment** [**terminate-only**]
no local-address-assignment

Context debug>service>id>ppp>event

Description This command enables debugging for local-address-assignment events.

The **no** form of this command disables debugging.

Parameters **terminate-only** — Enables debugging for local address assignment.

ppp

Syntax **ppp** [**terminate-only**]
no ppp

Context debug>service>id>ppp>event

Description This command enables PPP event debug.

The **no** form of this command disables debugging.

Parameters **terminate-only** — Enables debugging for local terminated PPP session

mac

Syntax [**no**] **mac** *ieee-address*

Context debug>service>id>ppp

Description This command displays PPP packets for a particular MAC address.

The **no** form of this command disables debugging.

msap

Syntax [**no**] **msap** *msap-id*

Context debug>service>id>ppp

Description This command enables debugging for specific PPP MSAPs.

The **no** form of this command disables debugging.

packet

Syntax [**no**] **packet**

Context debug>service>id>ppp

Description This command enables the PPP packet debug context.

The **no** form of this command disables debugging.

detail-level

Syntax **detail-level** {**low** | **medium** | **high**}
no detail-level

Context debug>service>id>ppp>packet

Description This command specify the detail level of PPP packet debug output.

The **no** form of this command disables debugging.

Parameters **low** | **medium** | **high** — Specifies the detail level of PPP packet debug output.

dhcp-client

Syntax [**no**] **dhcp-client**

Context debug>service>id>ppp>packet

Description This command enables packet debug output for DHCP client of the PPP session

The **no** form of this command disables debugging.

discovery

Syntax **discovery** [**padi**] [**pado**] [**padr**] [**pads**] [**padt**]
no discovery

Context debug>service>id>ppp>packet

Description This command enables PPP discovery packet debug output.

The **no** form of this command disables debugging.

Parameters [**padi**] [**pado**] [**padr**] [**pads**] [**padt**] — Enables the corresponding type of PPP discovery packet.

mode

Syntax **mode** {**dropped-only** | **ingr-and-dropped** | **egr-ingr-and-dropped**}
no mode

Context debug>service>id>ppp>packet

Description This command specifies PPP packet debug mode.

The **no** form of this command disables debugging.

Parameters **dropped-only** — Displays only dropped packets.

ingr-and-dropped — Displays only ingress packet and dropped packets.

egr-ingr-and-dropped — Displays ingress, egress and dropped packets.

ppp

Syntax **ppp** [**lcp**] [**pap**] [**chap**] [**ipcp**] [**ipv6cp**]
no ppp

Context debug>service>id>ppp>packet

Description This command enables PPP discovery packet debug output for the specified PPP protocol.

The **no** form of this command disables debugging.

Parameters [**lcp**] [**pap**] [**chap**] [**ipcp**] [**ipv6cp**] — Enables debug for the specified protocol.

remote-id

Syntax [**no**] **remote-id** *remote-id*

Context debug>service>id>ppp

Description This command enables debugging for specific PPP remote-ids.

The **no** form of this command disables debugging.

sap

Syntax [**no**] **sap** *sap-id*

Context debug>service>id>ppp

Description This command enables PPP debug output for the specified SAP, this command allow multiple instances.

The **no** form of this command disables debugging.

Parameters *sap-id* — Specifies the SAP ID.

username

Syntax [**no**] **username** *username*

Context debug>service>id>ppp

Description This command enable PPP debug for the specified username. since not all PPP packets contain username, so a MAC debug filter is created automatically when system sees a PPP packet contain the specified username.

Multiple username filters can be specified in the same debug command.

The **no** form of this command disables debugging.

Parameters *user-name* — Specifies the PPP username.

circuit-id

Syntax [**no**] **circuit-id** *circuit-id*

Context debug>service>id>ppp

Description This command enable PPP debug for the specified circuit-id.

Multiple circuit-id filters can be specified in the same debug command.

The **no** form of this command disables debugging.

Parameters *circuit-id* — Specifies the circuit-id in PADI.

remote-id

Syntax [**no**] **remote-id** *remote-id*

Context debug>service>id>ppp

Description This command enable PPP debug for the specified remote-id.

Multiple remote-id filters could be specified in the same debug command.

Parameters *remote-id* — Specifies the remote-id in PADI.

msap

Syntax [**no**] **msap** *msap-id*

Context debug>service>id>ppp

Description This command enable PPP debug for the specified managed SAP.

Multiple msap filters could be specified in the same debug command.

The **no** form of this command disables debugging.

Parameters *msap-id* — Specifies the managed SAP ID.

authentication

Syntax **authentication** [**policy** *policy-name*] [**mac-addr** *ieee-address*] [**circuit-id** *circuit-id*]

Context debug>subscr-mgmt

Description This command debugs subscriber authentication.

Parameters *policy-name* — Specifies an existing subscriber management authentication policy name.

ieee-address — Specifies the 48-bit MAC address xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx.

circuit-id — Specify the circuit-id, up to 256 characters.

sub-ident-policy

Syntax [**no**] **sub-ident-policy** *policy-name*

Context debug>subscr-mgmt

Description This command debugs subscriber identification policies.

The **no** form of this command disables debugging.

Parameters *policy-name* — Specifies the subscriber identification policy to debug.

script-compile-error

Syntax [**no**] **script-compile-error**

Context debug>subscr-mgmt>sub-ident-plcy

Description This command send the traceback of the compile error to the logger. The traceback contains detailed information about where and why the compilation fails. The compilation takes place when the CLI user changes the admin state of the Python URL from shutdown to no-shutdown.

The **no** form of this command disables debugging.

script-export-variables

Syntax [**no**] **script-export-variables**

Context debug>subscr-mgmt>sub-ident-plcy

Description This command sends the result (the three output variables) of the Python script to the logger when the script ran successfully.

The **no** form of this command disables debugging.

script-output

Syntax [**no**] **script-output**

Context debug>subscr-mgmt>sub-ident-plcy

Description This command sends the output (such as from 'print' statements) of the Python script to the logger.

The **no** form of this command disables debugging.

script-output-on-error

Syntax [**no**] **script-output-on-error**

Context debug>subscr-mgmt>sub-ident-plcy

Description This command sends the output (such as from print statements) of the Python script to the logger, but only when the script fails.

The **no** form of this command disables debugging.

script-runtime-error

Syntax [**no**] **script-runtime-error**

Context debug>subscr-mgmt>sub-ident-plcy

Description This command sends the traceback of the Python script failure to the logger.

The **no** form of this command disables debugging.

script-all-info

Syntax **script-all-info**

Context debug>subscr-mgmt>sub-ident-plcy

Description This command enables the script-compile-error, script-export-variables, script-output, script-output-on-error, and script-runtime-error functionalities.

srrp

Syntax [**no**] **srrp**

Context debug>router

Description This command enables debugging for SRRP packets.

The **no** form of this command disables debugging.

events

Syntax [**no**] **events** [**interface** *ip-int-name*]

Context debug>router>srrp

Description This command enables debugging for SRRP packets.

The **no** form of this command disables debugging.

packets

Syntax [**no**] **packets** [**interface** *ip-int-name*]

Context debug>router>srrp

Description This command enables debugging for SRRP packets.

The **no** form of this command disables debugging.

radius

Syntax [**no**] **radius**

Context debug>router

Description This command enables the debug router RADIUS context.

detail-level

Syntax **detail-level** {**low** | **medium** | **high**}
no detail-level

Context debug>router>radius

Description This command specifies the output detail level of command **debug router radius**.

Default detail-level medium

Parameters **low** — Specifies that the output includes packet type, server address, length, radius-server-policy name.

medium — Specifies all output in low level including the RADIUS attributes in the packet.

high — Specifies all output in medium level including the hex packet dump.

packet-type

Syntax **packet-type** [**authentication**] [**accounting**] [**coa**]
no packet-type

Context debug>router>radius

Description This command specifies the RADIUS packet type filter of command **debug router radius**

Default authentication accounting coa

Parameters **authentication** — Specifies the RADIUS authentication packet.

accounting — Specifies the RADIUS accounting packet.

coa — Specifies the RADIUS change of authorization packet.

radius-attr

Syntax **radius-attr type** *attribute-type* [**extended-type** *attribute-ext-type*] [**transaction**]
radius-attr type *attribute-type* [**transaction**] {**address** | **hex** | **integer** | **string**} **value** *attribute-value*
radius-attr vendor *vendor-id* **type** *attribute-type* [**extended-type** *attribute-ext-type*] [**transaction**] [**encoding** *encoding-type*]
radius-attr vendor *vendor-id* **type** *attribute-type* [**extended-type** *attribute-ext-type*] [**transaction**] [**encoding** *encoding-type*] {**address** | **hex** | **integer** | **string**} **value** *attribute-value*
no radius-attr type *attribute-type* [**extended-type** *attribute-ext-type*]
no radius-attr type *attribute-type* [**extended-type** *attribute-ext-type*] {**address** | **hex** | **integer** | **string**} **value** *attribute-value*
no radius-attr vendor *vendor-id* **type** *attribute-type* [**extended-type** *attribute-ext-type*]
no radius-attr vendor *vendor-id* **type** *attribute-type* [**extended-type** *attribute-ext-type*] {**address** | **hex** | **integer** | **string**} [0..16777215] *attribute-value*

Context debug>router>radius

Description This command specifies the RADIUS attribute filter of command **debug router radius**.

Parameters *attribute-type* — Specifies the RADIUS attribute type.

Values 1 to 255

attribute-ext-type — Specifies the RADIUS attribute extended type (RFC 6929).

Values 1 to 255

address — Specifies the value is a IPv4 or IPv6 address/prefix/subnet.

string — Specifies the value is a ASCII string.

integer — Specifies the value is a integer.

hex — Specifies the value is a binary string in hex format, such as “\0xAB01FE”.

attribute-value — Specifies the value of the RADIUS attribute.

Values

address	<ipv4-address> | <ipv6-address> | <ipv6-prefix/prefix-length>	
	ipv4-address	a.b.c.d

ipv6-address	x:x:x:x:x:x:x:x (eight 16-bit pieces)
ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d	
x - [0 to FFFF]H	
d - [0 to 255]D	
ipv6-prefix-length [0 to 128]	
hex	[0x0 to 0xFFFFFFFF (max 506 hex nibbles)]
integer	[0 to 4294967295]
string	ascii-string (max 253 chars)

transaction — Specifies that the system outputs both request and response packets in the same session even if the response packet does not include the filter attribute.

vendor-id — Specifies the vendor id for the vendor specific attribute.

Values 0 to 16777215

encoding-type — Specifies the size of the vendor-type and vendor-length in bytes. It is a two digitals string: “xy”, x is the size of vendor-type, range from 1 to 4; y is the size of vendor-length, range from 0 to 2; it is “11” by default.

Values type-size:1 to 4, length-size: 0 to 2

wpp

Syntax [**no**] **wpp**

Context debug>router

Description This command enables the context to configure WPP debugging parameters.

packet

Syntax [**no**] **packet**

Context debug>router>wpp
debug>router>wpp>portal

Description This command enables WPP packet debugging.

detail-level

Syntax **detail-level** *detail-level*

Context debug>router>wpp
debug>router>wpp>packet
debug>router>wpp>portal>packet

Description This command specifies the detail level of WPP packet debugging.

Parameters *detail-level* — specifies the detail level of WPP packet debugging.

Values high — Specifies a high detail level for WPP packet debugging
low — Specifies a low detail for WPP packet debugging

portal

Syntax [**no**] **portal** *wpp-portal-name*

Context debug>router>wpp

Description This command enables WPP debugging for the specified WPP portal.

Parameters *wpp-portal-name* — Specifies the WPP portal name.

call-trace

Syntax **call-trace**

Context debug

Description This command enables the context to set up various call trace debug sessions.

ipoe

Syntax **ipoe**

Context debug>call-trace

Description This command enables the context to set up call trace debugging for IP over Ethernet (IPoE) sessions.

trace

Syntax **trace sap** *sap-id* [**mac** *ieee-address*] [**circuit-id** *circuit-id*] [**remote-id** *remote-id*] [**profile** *trace-profile-name*] [**trace-existing-sessions**] [**max-jobs** *num*] [**name** *trace-name*]
trace mac *ieee-address* [**sap** *sap-id*] [**circuit-id** *circuit-id*] [**remote-id** *remote-id*] [**profile** *trace-profile-name*] [**trace-existing-sessions**] [**max-jobs** *num*] [**name** *trace-name*]
trace circuit-id *circuit-id* [**profile** *trace-profile-name*] [**trace-existing-sessions**] [**max-jobs**

num] [**name** *trace-name*]

trace remote-id *remote-id* [**profile** *trace-profile-name*] [**trace-existing-sessions**] [**max-jobs** *num*] [**name** *trace-name*]

no trace [**sap** *sap-id*] [**mac** *ieee-address*] [**circuit-id** *circuit-id*] [**remote-id** *remote-id*]

no trace name *trace-name*

Context debug>call-trace>ipoe

Description This command enables tracing for IPoE sessions specified by the configured parameters. This command can trace a single session or multiple sessions, and can use wildcard characters.

This command can be executed multiple times to start multiple traces. When rules overlap, such as for a wildcard SAP and a specific SAP, the rule that a specific trace is associated with cannot be guaranteed.

The **no** form of this command prevents new traces from being configured and terminates all trace jobs that were previously started using the **trace** command.

Parameters *circuit-id* — Specifies a circuit ID that is used to filter sessions to trace. The *circuit-id* and *remote-id* parameters are mutually exclusive.

ieee-address — Specifies a MAC address that is used to identify a session to trace, in the format “ab:cd:ef:01:23:45”. A wildcard character can be used to match all remaining octets; for example, the format “ab:cd:ef:*” can be used to filter by OUI.

num — Specifies the maximum number of jobs that may be started with this rule.

Values 1 to 50

Default 1

remote-id — Specifies a remote ID that is used to filter sessions to trace. The *remote-id* and *circuit-id* parameters are mutually exclusive.

sap-id — Specifies a SAP to trace. The following formats are accepted:

- *port*/*lag*/*pw-port*:*svlan.cvlan*
- *port*/*lag*/*pw-port*:*vlan*
- *port*/*lag*/*pw-port*
- *port*/*lag*/*pw-port*:*vlan*.*
- *port*/*lag*/*pw-port*:* (also matches *.*)

trace-existing-sessions — Specifies that existing IPoE sessions is traced. If this parameter is not included, only new IPoE sessions is traced.

trace-name — Specifies the name by which the trace is referenced, up to 16 characters maximum.

trace-profile-name — Specifies the name of the trace profile to be applied. The default parameters is used if a trace profile is not specified.

9.25.2.6 Monitor Commands

subscriber

Syntax **subscriber** *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**] [**base**] [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]

subscriber *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**] **egress-policer-id** *egress-policer-id* [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]

subscriber *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**] **egress-queue-id** *egress-queue-id* [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]

subscriber *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**] **ingress-policer-id** *ingress-policer-id* [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]

subscriber *sub-ident-string* **sap** *sap-id* **sla-profile** *sla-profile-name* [**interval** *seconds*] [**repeat** *repeat*] [**absolute** | **rate**] **ingress-queue-id** *ingress-queue-id* [**spi-sharing-type** *spi-sharing-type*] [**spi-sharing-id** *spi-sharing-id*]

Context monitor>service

Description This command monitors statistics for a subscriber.

Parameters *sub-ident-string* — Specifies an existing subscriber identification profile to monitor.

sap-id — Specifies the physical port identifier portion of the SAP definition.

sla-profile-name — Specifies an existing SLA profile.

seconds — Configures the interval for each display in seconds.

Default 11

Values 11 to 60

repeat — Configures how many times the command is repeated.

Default 10

Values 1 to 999

absolute — Specifies that raw statistics are displayed, without processing. No calculations are performed on the delta or rate statistics.

Default mode delta

rate — Specifies that the rate-per-second for each statistic is displayed instead of the delta.

base — Monitors base statistics.

egress-policer-id — Monitors statistics for this egress policer.

Values 1 to 63

egress-queue-id — Monitors statistics for this egress queue.

Values 1 to 8

ingress-policer-id — Monitors statistics for this ingress policer.

Values 1 to 63

ingress-queue-id — Monitors statistics for this ingress queue.

Values 1 to 32

spi-sharing-type — Monitors statistics for sessions with specified SPI sharing type.

Values per-group, per-session-ipoe, per-session-ppp

spi-sharing-id — Monitors statistics for sessions with specified SPI sharing type and sharing identifier. The SPI sharing identifier is:

- an identifier when the sharing type is per-group
- an internal PPP or IPoE identifier when the sharing type is per-PPP or per-IPoE session

Values 0 to 4294967295

Output The following is an example of subscriber monitor statistics.

Sample Output

```
A:Dut-A# monitor service subscriber nokia_100 sap 1/2/1:101 sla-profile sla_default
===============================================================================
Monitor statistics for Subscriber nokia_100
===============================================================================
At time t = 0 sec (Base Statistics)
-------------------------------------------------------------------------------
SLA Profile Instance statistics
-------------------------------------------------------------------------------
                          Packets                 Octets
Off. HiPrio           : 0                         0
Off. LowPrio          : 94531                   30704535
Off. Uncolor          : 0                         0

Queueing Stats (Ingress QoS Policy 1000)
Dro. HiPrio           : 0                         0
Dro. LowPrio          : 7332                      2510859
For. InProf           : 0                         0
For. OutProf          : 87067                     28152288

Queueing Stats (Egress QoS Policy 1000)
Dro. InProf           : 880                       127660
Dro. OutProf          : 0                         0
For. InProf           : 90862                     12995616
For. OutProf          : 0                         0
-------------------------------------------------------------------------------
SLA Profile Instance per Queue statistics
-------------------------------------------------------------------------------
```

```
                        Packets                 Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Off. Uncolor          : 0                       0
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Queue 2 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LowPrio          : 94531                   30704535
Off. Uncolor          : 0                       0
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 7332                    2510859
For. InProf           : 0                       0
For. OutProf          : 87067                   28152288

Ingress Queue 3 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Off. Uncolor          : 0                       0
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Ingress Queue 11 (Multipoint) (Priority)
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
Off. Uncolor          : 0                       0
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Egress Queue 1
Dro. InProf           : 880                     127660
Dro. OutProf          : 0                       0
For. InProf           : 90862                   12995616
For. OutProf          : 0                       0

Egress Queue 2
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0

Egress Queue 3
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
===============================================================================
A:Dut-A#


A:Dut-A# monitor service subscriber nokia_100 sap 1/2/1:101 sla-
```

```
profile sla_default base rate
===============================================================================
Monitor statistics for Subscriber nokia_100
===============================================================================
At time t = 0 sec (Base Statistics)
-------------------------------------------------------------------------------
SLA Profile Instance statistics
-------------------------------------------------------------------------------
                        Packets                 Octets
Off. HiPrio           : 0                       0
Off. LowPrio          : 109099                  35427060
Off. Uncolor          : 0                       0

Queueing Stats (Ingress QoS Policy 1000)
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 8449                    2894798
For. InProf           : 0                       0
For. OutProf          : 100523                  32489663

Queueing Stats (Egress QoS Policy 1000)
Dro. InProf           : 880                     127660
Dro. OutProf          : 0                       0
For. InProf           : 105578                  15104553
For. OutProf          : 0                       0
-------------------------------------------------------------------------------
At time t = 11 sec (Mode: Rate)
-------------------------------------------------------------------------------
SLA Profile Instance statistics
-------------------------------------------------------------------------------
                        Packets                 Octets                  % Port
                                                                        Util.
Off. HiPrio           : 0                       0                       0.00
Off. LowPrio          : 1469                    477795                  0.38
Off. Uncolor          : 0                       0                       0.00

Queueing Stats (Ingress QoS Policy 1000)
Dro. HiPrio           : 0                       0                       0.00
Dro. LowPrio          : 119                     40691                   0.03
For. InProf           : 0                       0                       0.00
For. OutProf          : 1349                    437350                  0.34

Queueing Stats (Egress QoS Policy 1000)
Dro. InProf           : 0                       0                       0.00
Dro. OutProf          : 0                       0                       0.00
For. InProf           : 1469                    209129                  0.16
For. OutProf          : 0                       0                       0.00
===============================================================================
A:Dut-A#
A:Dut-A# monitor service subscriber nokia_100 sap 1/2/1:101 sla-
profile sla_default ingress-queue-id 1
===============================================================================
Monitor statistics for Subscriber nokia_100
===============================================================================
At time t = 0 sec (Base Statistics)
-------------------------------------------------------------------------------
                        Packets                 Octets
Ingress Queue 1 (Unicast) (Priority)
Off. HiPrio           : 0                       0
Off. LowPrio          : 0                       0
```

```
Off. Uncolor          : 0                       0
Dro. HiPrio           : 0                       0
Dro. LowPrio          : 0                       0
For. InProf           : 0                       0
For. OutProf          : 0                       0
===============================================================================
A:Dut-A#


A:Dut-A# monitor service subscriber nokia_100 sap 1/2/1:101 sla-
profile sla_default egress-queue-id 1
===============================================================================
Monitor statistics for Subscriber nokia_100
===============================================================================
At time t = 0 sec (Base Statistics)
-------------------------------------------------------------------------------
                        Packets                 Octets
Egress Queue 1
Dro. InProf           : 880                     127660
Dro. OutProf          : 0                       0
For. InProf           : 164366                  23506178
For. OutProf          : 0                       0
===============================================================================
A:Dut-A#
```

host

Syntax **host** [**sap** *sap-id*] [**wholesaler** *service-id*] [**port** *port-id*] [**inter-dest-id** *intermediate-destination-id*] [**detail**]

host [**sap** *sap-id*] [**wholesaler** *service-id*] [**port** *port-id*] **no-inter-dest-id** [**detail**]

host summary

host [**detail**] **wholesaler** *service-id* (**VPRN only**)

Context show>service>id

Description This command displays static host information configured on this service.

Parameters *sap-id* — Displays SAP information for the specified SAP ID.

intermediate-destination-id — Specifies the intermediate destination identifier which is encoded in the identification strings up to 32 characters.

summary — Displays summary host information.

wholesaler *service-id* — Specifies the VPRN service ID of the wholesaler. When specified in this context, SAP, SDP, interface, IP address and MAC parameters are ignored. This parameter applies only to the 7750 SR.

Values

service-id:	1 to 2147483647
svc-name:	64 characters maximum

10 Oversubscribed Multi-Chassis Redundancy (OMCR) in ESM

10.1 Overview

10.1.1 Terminology and Abbreviations

- **OMCR** — Oversubscribed Multi-Chassis Redundancy
- **Warm-Standby Node** or **Protecting Node** — Refers to the oversubscribed node that offers the protection of subscriber hosts spread over multiple BNGs. During the normal operation, the protecting node maintains the subscriber host in the form of an MCS record (Multi-Chassis Synchronization Record) in the control plane. Only when the failure occurs and the protecting node becomes active, are the subscriber-hosts fully instantiated in the data and control plane. This node is sometimes referred to as N:1 node.
- **Active/Active (1:1) Model** — This mode of operation refers to the model where subscribers host are fully synchronized between two chassis, regardless of the state of the underlying SRRP instance (Master/Standby). Each node can have MCS peering sessions with four another nodes where each peering session represent 1 to 1 mapping set of active subscriber hosts.

10.1.2 Restrictions

- The protecting node must use CPM-4 or higher (other protected nodes can continue to use CPM-3).
- The protecting node functionality is not supported in mixed-mode in 7450 ESS chassis.
- All nodes in the OMCR cluster (central standby and the protected nodes) must run at the minimum SR OS Release 12.0R1.
- Warm-standby mode is a chassis-wide property. In other words, while in warm-standby mode, the chassis cannot operate in 1:1 (active-active) redundancy mode.

- OMCR is supported only for IPoEv4/v6 subscribers. However, non-synchronized PPPoEv4/v6 subscribers are supported in the OMCR cluster. PPPoEv4/v6 PTA (locally terminated) non-synchronized subscribers and the OMCR synchronized IPoE subscriber must be instantiated under separate group-interfaces. On the other hand, non-synchronized PPPoEv4/v6 LAC sessions are allowed to be under the same group interface as the OMCR synchronized IPoE subscribers. Non-synchronized PPPoEv4/v6 subscriber hosts will rely on ppp-keepalive timeouts to re-establish connectivity when the failure occurs.
- Preemption of already instantiated subscriber hosts in the protecting node by another subscriber hosts is not allowed.
- Redundant interface (shunting) is not supported for subscribers on the protecting node while they are not fully instantiated in the control/data plane (or while the underlying SRRP instance is in a non-Master state on the protecting node).
- Persistency in multi-chassis environment **must** be disabled since redundant nodes are protecting each other and they maintain up-to-date lease states.
- The failover trigger is based on SRRP only (no MC-LAG support).
- Unnumbered subscriber-interface model is not supported in OMCR.
- The protecting node supports 10 MCS peers, while the protected node (in active/active mode of operation) supports 4 MCS peers.
- Synchronization of the following MCS clients is not supported:
  - Host tracking
  - MC ring
  - Layer 2 subscriber hosts
  - Layer 3 IGMP/MLD
  - Layer 2 IGMP/MLD
  - DHCP Server
  - PPPoE Clients
  - MC-LAG
  - MC-IPSEC
  - MC-ENDPOINT

10.2 Deploying Oversubscribed Multi-Chassis Redundancy

In order to optimize the cost, certain operators prefer oversubscribed model in which a single central standby BNG (protecting BNG) supports multiple other BNGs in a semi-stateful fashion.

In Oversubscribed Multi-Chassis Redundancy (OMCR) model, a large number of subscriber-hosts are backed up by a single central standby node. Standby subscriber-hosts within the protecting node are synchronized only within the control plane (CPM) in the form of a Multi-Chassis Synchronization (MCS) record. Such subscriber hosts are not instantiated in the data plane and therefore the data plane resources can be spared and used only on an **as needed** basis. This trait allows the protecting node to back up a large number of subscribers that are scattered over multiple active BNG nodes at the expense of slower convergence.

Only a subset of the subscribers, up to the available resource capacity of the data plane in the protecting node, would be activated on the protecting node at any given time during the failure.

The failover trigger is based on SRRP (no MC-LAG support). The subscriber hosts under the corresponding group-interface is switched over once the SRRP instance on the protecting node transitions into the Master SRRP state.

There are two possible models for this deployment:

1. Access nodes are directly connected to the BNGs. From the perspective of standby subscribers, in this model the line card is oversubscribed but the physical ports on it are not. For example each of the 10 physical ports on the same line card can be directly connected to respective access nodes. Assume that each physical port can support 64K subscriber hosts. Considering that the subscriber host limit per line card is also 64K (at the time of this writing), the oversubscription ratio in this case would be 10:1.

   The concept of this deployment scenario is shown in Figure 170.

Figure 170 **OMCR Scenario Without Aggregation Network**

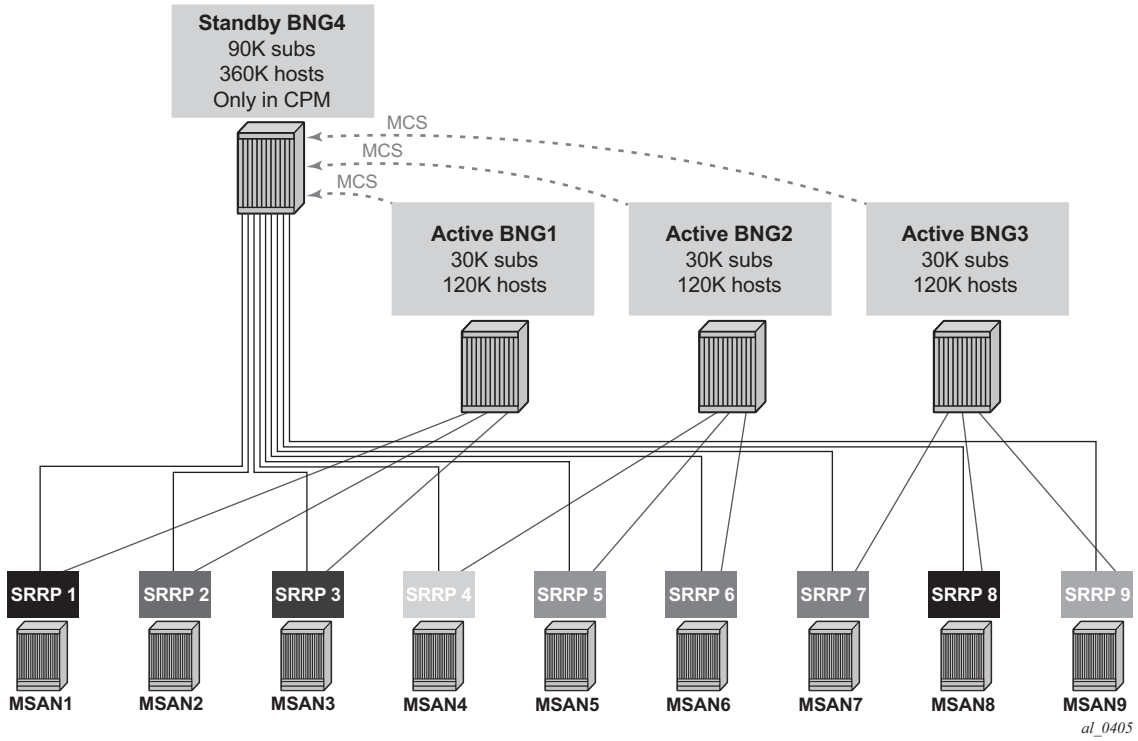

2. Aggregation network in the access (double VLAN tags). In this case a line card and a physical port can be oversubscribed with standby subscribers. For example multiple capture-saps (each capture sap containing 4K c-vlans) can be created on a single physical port on the protecting BNG, for the total of >>64K subscribers per physical port.

   Conceptual model for this scenario is shown in Figure 171 (although the number of SRRP instances and capture-saps is in this figure is reduced for simplification).

Figure 171 **OMCR Scenario with Aggregation Network**

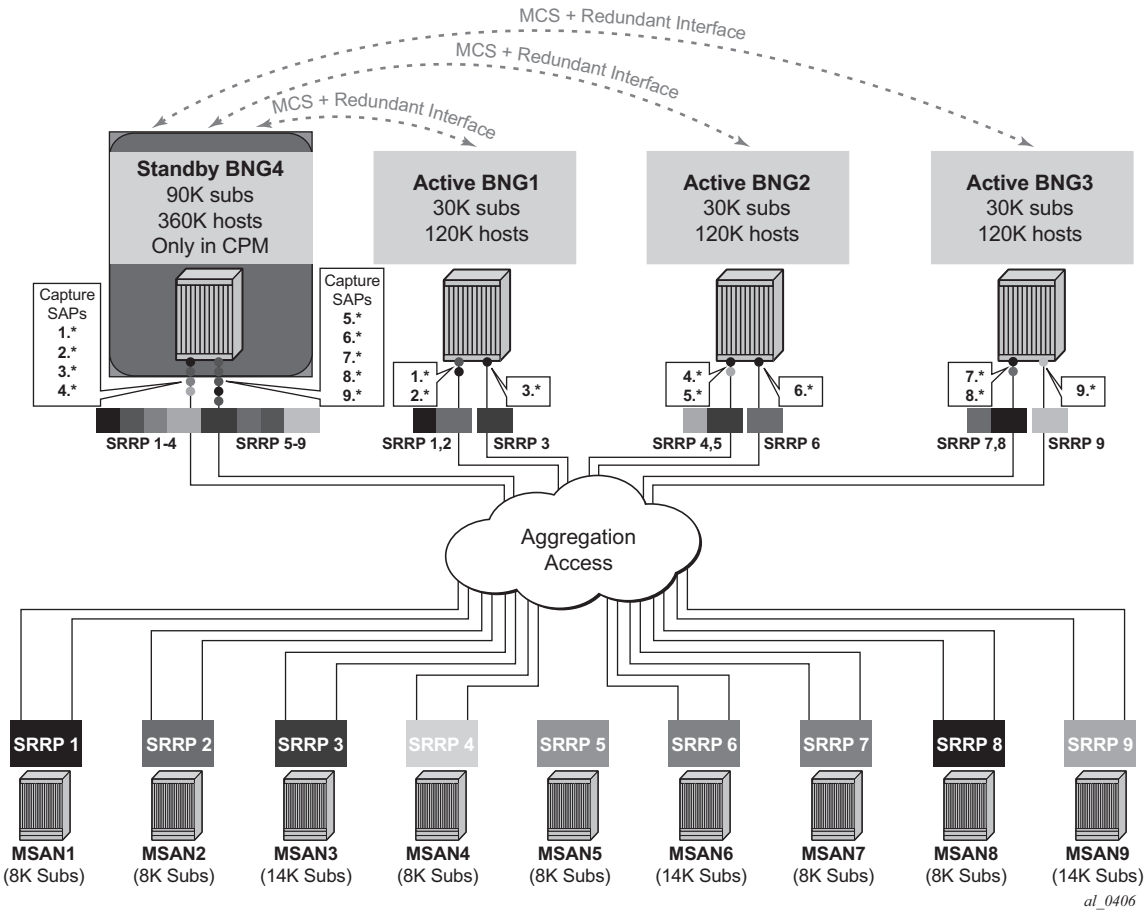

In both cases, a maximum of 64K subscribers per line card can be activated on the protecting BNG during the switchover. This is something that the operator should plan around, and consequently group the access nodes in a way so that the eventual number of active subscribers per line card on the protecting node does not exceed the maximum number of supported subscribers per line card.

Note: One could have deployment scenario in which system wide ESM capacity is oversubscribed but the line card capacity is not. For example, on chassis with 10 line cards, each line card can be reserved to protect a total host count of 64k. This would yield a total of 640k protected hosts distributed across the 10 cards but only up to 256k hosts could be activated simultaneously should it be required due to SRRP transitions to Master.

10.2.1 Resource Exhaustion Notification and Simultaneous Failures

The protection success of the OMCR model relies on grouping protected entities (links and nodes) according to the likelihood of their failure within the time frame required for their restoration. For example the same resource (IOM card or port) on the protecting node can be used to protect multiple entities in the network as long as their failures do not overlap in time. In other words, if one failure can be repaired before the next one contending for the same resource on the central standby node, the OMCR model will serve the purpose.

But since the oversubscribed model does not offer any guarantees, it is possible that the protecting node in certain cases runs out of resources and fails to offer protection. In this case, the protecting node will generate a SNMP trap identifying the SRRP instance on which subscriber protection has failed. One SNMP trap is raised per SRRP instance in case where at least one subscriber under the corresponding group interface was not instantiated. The trap is cleared either when all subscribers become instantiated or when the SRRP transition into a non-master state.

The number of the subscriber hosts that failed to instantiate, can also be determined using the operational **show>redundancy>multi-chassis all** command. This command will show the number of subscribers that failed to instantiate along with SRRP instances on which the subscriber host are relaying for successful connectivity.

Pre-emption of already instantiated subscriber hosts in the protecting node by another subscriber hosts is not allowed.

10.2.2 Resource Monitoring

Management and conservation of resources is of utmost importance in OMCR. The resources consumed by the subscriber host depend on the type and the size of subscriber parameters (the number of strings, length of strings, and so on).

For these reasons it is crucial that the operator has a view of the amount of memory in the CPM utilized by subscribers and the amount of free memory that can be used for additional subscribers. The **MCS** line is of particular interest in this output. In addition, the **Subscriber Mgmt** line shows memory utilization for active subscribers in the CPM.

The **Available Memory** gives an indication about how much memory remains.

For example:

```
*A:right-21# show system memory-pools
===============================================================================
Memory Pools
===============================================================================
Name                Max Allowed    Current Size      Max So Far          In Use
-------------------------------------------------------------------------------
BFD                    No limit       6,291,456       6,291,456       5,509,872
BGP                    No limit       5,242,880       5,242,880       3,635,976
CFLOWD                 No limit       1,048,576       1,048,576          26,576
Cards & Ports          No limit      24,117,248      27,262,976      18,010,424
DHCP Server            No limit       2,097,152       2,097,152         173,680
ETH-CFM                No limit       6,291,456       9,437,184       4,016,128
ICC                 25,165,824       7,340,032      25,165,824       2,880,008
IGMP/MLD               No limit       1,048,576       1,048,576         166,216
IMSI Db Appl           No limit       1,048,576       1,048,576         793,984
IOM                    No limit       8,388,608       8,388,608       6,894,360
IP Stack               No limit      29,360,128      35,651,584      13,565,120
IS-IS                  No limit       2,097,152       2,097,152       1,095,360
ISA                    No limit       3,145,728       3,145,728       1,217,464
LDP                    No limit       6,291,456       6,291,456       5,607,240
Logging             411,041,792       6,291,456       6,291,456       3,473,024
MBUF              1,073,741,824       2,097,152       2,097,152         299,976
MCS                    No limit     454,033,408     454,033,408     416,753,472
MPLS/RSVP              No limit      49,283,072      69,206,016      42,947,776
MSCP                   No limit       2,097,152       2,097,152       1,022,848
MSDP                   No limit               0               0               0
Management             No limit      19,922,944      26,214,400       5,689,112
OAM                    No limit       1,048,576       1,048,576          86,080
OSPF                   No limit       8,388,608       8,388,608       4,975,824
OpenFlow               No limit       1,048,576       1,048,576         391,880
PIM                    No limit      19,922,944      19,922,944      15,755,792
PTP                    No limit       1,048,576       1,048,576           1,408
RIP                    No limit               0               0               0
RTM/Policies           No limit       9,437,184       9,437,184       7,002,648
Redundancy             No limit       9,437,184     424,673,280         703,160
SIM                    No limit       3,145,728      12,582,912             648
Services               No limit      25,165,824      25,165,824      18,128,056
Stats                  No limit       1,048,576       1,048,576           9,456
Subscriber Mgmt        No limit      24,117,248      41,943,040      14,846,512
System                 No limit     794,820,608     856,686,592     776,394,656
Traffic Eng            No limit       1,048,576       1,048,576         444,744
VRRP                   No limit       2,097,152       3,145,728         393,808
WEB Redirect        16,777,216       1,048,576       1,048,576         128,640
-------------------------------------------------------------------------------
Current Total Size :    1,540,358,144 bytes
Total In Use       :    1,373,041,928 bytes
Available Memory   :    5,778,702,336 bytes
===============================================================================
```

Similar output is given in regards to CPU utilization:

```
*A:right-21#  show system cpu
===============================================================================
CPU Utilization (Sample period: 1 second)
===============================================================================
Name                                    CPU Time       CPU Usage        Capacity
                                          (uSec)                          Usage
-------------------------------------------------------------------------------
```

```
BFD                                          0          0.00%          0.00%
BGP                                      2,504          0.03%          0.05%
BGP PE-CE                                    0          0.00%          0.00%
CFLOWD                                      25         ~0.00%         ~0.00%
Cards & Ports                           14,501          0.18%          0.13%
DHCP Server                                 30         ~0.00%         ~0.00%
ETH-CFM                                    614         ~0.00%          0.06%
ICC                                      1,803          0.02%          0.18%
IGMP/MLD                                   538         ~0.00%          0.05%
IMSI Db Appl                                37         ~0.00%         ~0.00%
IOM                                          0          0.00%          0.00%
IP Stack                                 4,578          0.05%          0.24%
IS-IS                                      423         ~0.00%          0.02%
ISA                                      2,690          0.03%          0.10%
LDP                                         78         ~0.00%         ~0.00%
Logging                                     13         ~0.00%         ~0.00%
MBUF                                         0          0.00%          0.00%
MCS                                      2,718          0.03%          0.27%
MPLS/RSVP                                1,137          0.01%          0.08%
MSCP                                         0          0.00%          0.00%
MSDP                                         0          0.00%          0.00%
Management                               6,571          0.08%          0.19%
OAM                                      1,532          0.01%          0.09%
OSPF                                    18,397          0.23%          0.08%
OpenFlow                                    18         ~0.00%         ~0.00%
PIM                                          0          0.00%          0.00%
PTP                                         24         ~0.00%         ~0.00%
RIP                                          0          0.00%          0.00%
RTM/Policies                                 0          0.00%          0.00%
Redundancy                               3,618          0.04%          0.19%
SIM                                     10,959          0.13%          1.08%
SNMP Daemon                                  0          0.00%          0.00%
Services                                 1,037          0.01%          0.03%
Stats                                        0          0.00%          0.00%
Subscriber Mgmt                            835          0.01%          0.03%
System                                  29,863          0.37%          1.32%
Traffic Eng                                  0          0.00%          0.00%
VRRP                                       970          0.01%          0.07%
WEB Redirect                                26         ~0.00%         ~0.00%
-------------------------------------------------------------------------------
Total                                7,975,383        100.00%
   Idle                              7,869,844         98.67%
   Usage                               105,539          1.32%
Busiest Core Utilization                33,264          3.33%
===============================================================================
*A:right-21#
```

10.2.3 Warm-Standby Mode Of Operation

The protecting node operates in a warm-standby mode. Warm-standby mode of operation is a property of the entire node. In other words, while in the central-standby mode of operation (warm-standby command), only subscribers under the SRRP instances that are in the Master state is fully instantiated in the data plane on the central standby node (protecting node). All other subscribers (under the SRRP

instances that are in the standby state) is synchronized only in the control plane. However, non-central standby nodes can have a peering connection with a protecting node (OMCR) and at the same time another peering connection with another active BNG node in active/active model. All nodes participating in the OMCR mode of operation must run SR OS Release 12.0 or higher. This model is shown in Figure 172.

Figure 172 **Network Wide Mixing of OMCR and Active/Active (1:1) Model**

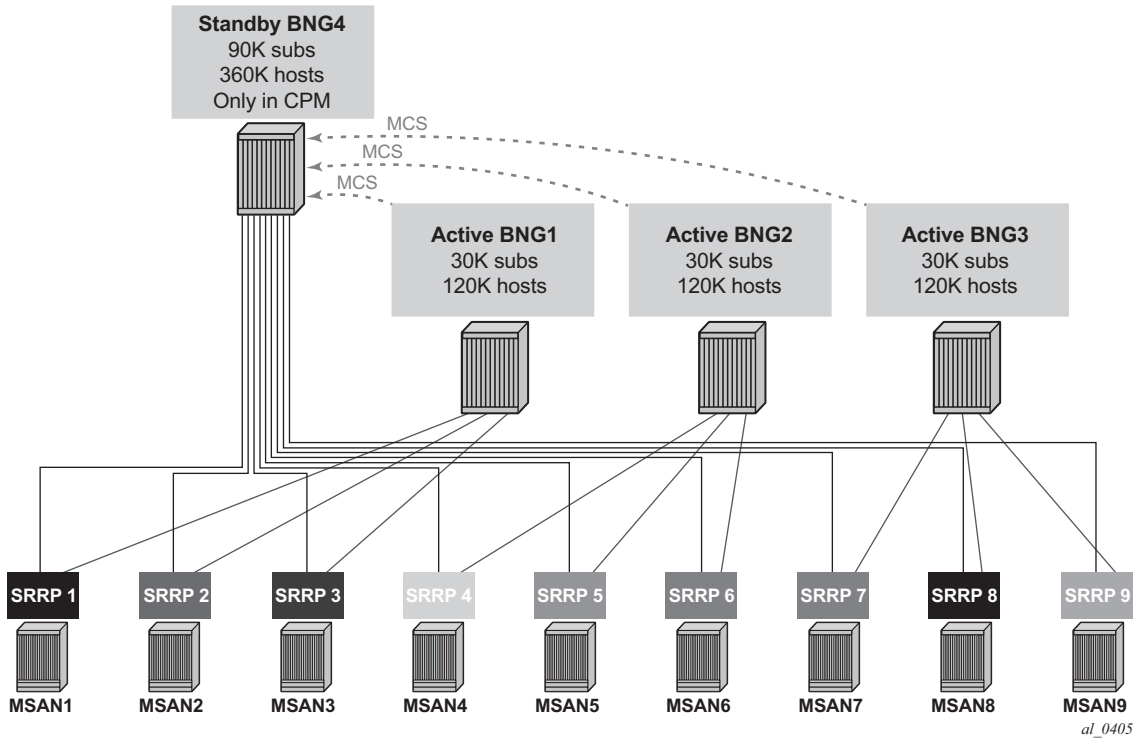

The central backup property is configured with the following CLI:

```
configure
    redundancy
        multi-chassis
            peer 10.1.1.1
                warm-standby
```

The **warm-standby** keyword configures the chassis to be in the central standby mode of operation. Although the configuration option is configured per peer, the **warm-standby** functionality is applied per chassis.

Synchronization of IPoE subscribers (**config>redundancy>multi-chassis>peer>sync>sub-mgmt ipoe**) on the protecting node is only possible if all peers are configured for **warm-standby** or none are.

To transition from one mode to another (warm — hot), all peers must be administratively shutdown and the warm-standby keyword must be either removed or configured on all peers, depending on the direction of the transition.

Single-homed subscribers are supported in the central standby node, subject to resource limitations.

10.2.4 IPoE vs PPPoE

OMCR is supported only for IPoEv4/v6 subscribers. PPPoEv4/v6 subscriber hosts are not supported. However, non-synchronized PPPoE hosts can be hosted on the protecting node simultaneously with the protected IPoE subscribers. PPPoE PTA (locally terminated) non-synchronized subscribers and OMCR synchronized IPoE subscriber must not be configured under the same group-interfaces. On the other hand, non-synchronized PPPoE LAC sessions are allowed to be under the same group interface as the OMCR synchronized IPoE subscribers.

The recovery of PPPoE subscriber host in non-synchronized environment is based on the timeout of ppp-keepalives.

10.2.5 Persistency

Persistency in the multi-chassis environment **must** be disabled since redundant nodes are protecting each other and they maintain up-to-date lease states. Otherwise, race conditions resulting in stale lease states caused by contention between MCS data and persistency data may occur.

10.2.6 Routing and Redundant Interface in OMCR

Support for redundant interface is limited and can be used only in cases where subscribers are activated in the protecting node. In other words, the shunting over the redundant interface cannot be used if subscriber hosts are not fully instantiated (in the data and control plane). For this reason, downstream traffic must not be attracted (via routing) to the protecting node while the subscriber hosts are in the standby mode (SRRP is in a backup state).

During the transient period while the switchover is in progress, subscriber hosts are being instantiated or withdrawn (depending on the direction of the switchover) in the data plane on the protecting node. The duration of this process is dependent on the number of the hosts that needs to be instantiated/withdrawn and it is proportional to the regular host setup/tear-down rates. The redundant interface in this case can be used only for the hosts that are present in the data plane during the switchover transitioning period (from the moment that they are instantiated in the dataplane when protecting node is assuming activity, or up the moment when they are withdrawn from the data plane when the protecting node is relinquishing activity).

The following routing models are supported:

Case 1 — SRRP-Aware routing where subnets can be assigned per group-interfaces (SRRP instances). In a steady state, the redundant interface is not needed since the downstream traffic is attracted to the master node. During switchover periods (routing convergence transitioning periods), redundant interface can be used only for the subscriber hosts that are instantiated in the data plane (from the moment that they are instantiated in the dataplane when protecting node is assuming activity, or up the moment when they are withdrawn from the data plane when the protecting node is relinquishing activity). See Figure 173.

Case 2 — SRRP-Aware routing where subnets spawn (are aggregated) over multiple group-interfaces (SRRP instances). In case of a switchover, /32 IPv4 addresses and /64 IPv6 addresses/prefixes are advertised from the protecting node. In a steady state, the redundant interface is not needed since the downstream traffic is attracted via more specific routes (/32s and /64s) to the master node. During switchover periods (routing convergence transitioning periods), the redundant interface can be used only for the subscriber-hosts that are instantiated in the data plane (from the moment that they are instantiated in the dataplane when protecting node is assuming activity, or up the moment when they are withdrawn from the data plane when the protecting node is relinquishing activity). To reduce the number of routes on the network side, /32s and /64s should only be activated on the protecting node.

Figure 173 **Subnet per Group Interface**

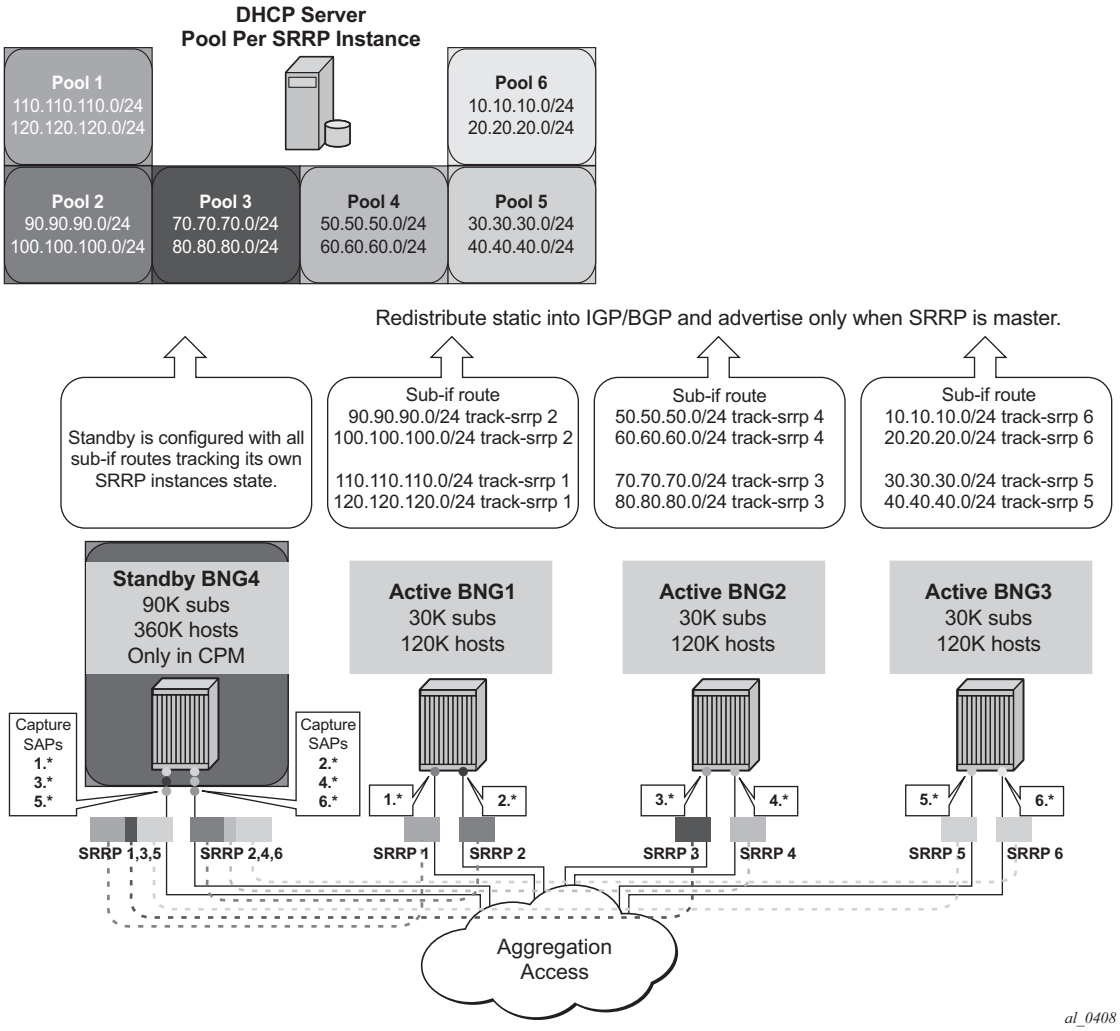

A deployment case that is not supported is the one where subnets spawn (are aggregated) over multiple group-interfaces (SRRP instances) and at the same time /32s are not allowed to be advertised from the protecting node. This scenario would require redundant interface support while subscriber-hosts are not necessarily instantiated in the protecting node.

10.2.7 Revertive Behavior

In case that failure is repaired on the original active node (non-central standby node) while SRRP preemption (**preempt**) is configured, the corresponding active subscribers on the protecting node is withdrawn from the data plane and the activity (mastership) is switched to the original node.

This behavior will ensure that the resources in the central backup are freed upon failure restoration and are available for protection of other entities in the network (other links/nodes).

In the preemption case, the upstream traffic is steered towards the newly active BNG via gratuitous ARP (GARP). In other words, the virtual MAC is advertised from the newly active node, and consequently the access and aggregation nodes will update their Layer 2 forwarding entries. This action should cause NO interruption in the upstream traffic.

In the downstream direction, the service interruption is equivalent to the time it takes to withdraw the routes from the network side on the standby node. In this case, there are two scenarios:

- A route per group interface (SRRP) is advertised in the network from the central standby node. In this case, downstream traffic interruption is a function of the convergence time of the routing protocol deployed on the network side. Once the routing is converged, all downstream subscriber traffic is attracted to the newly active node. In the meantime, the redundant interface can be used to shunt traffic from the central standby node to the newly master, but only for the subscriber hosts that have not yet been withdrawn from the data plane on the protecting node. This withdrawal process may take some time and therefore downstream traffic for some subscriber hosts is restored before the others during the routing convergence period.
- /32 subscriber-host routes are advertised from the protecting node. The total recovery time for downstream traffic will depend on the routing convergence. The routing convergence might be slower than in the previous case since more routes (/32s) need to be withdrawn from the network. The redundant interface can be used in the meantime for the subscriber hosts that have not yet been withdrawn from the data plane in protecting node.

10.2.8 Service Restoration Times

Service restoration times depends on the scale of the outage. The factors that affect the restoration times are:

- Failure detection time based on SRRP (could be in a sub second range, also supported based on BFD).
- Time needed to instantiate/withdraw subscriber host in/from the data plane.
- Routing convergence (based on SRRP aware routing).

10.2.9 Processing of the SRRP Flaps

When multiple srrp instances fail at the same time, they is processed one at the time on first come first serve basis. The subscriber instantiation processing during the switchover is divided into 1seconds intervals. In-between those intervals, the state of the SRRP are checked to ensure that it has not changed while the subscriber instantiation is in progress. This mechanism will break the inertia (snowball effect) that can be caused by SRRP instance flaps. Furthermore, an SRRP flap is handled by not requesting a withdrawal followed by an instantiation request for the same SRRP instance.

10.2.10 Accounting

The OMCR accounting follows the active/active (1:1) redundancy model.

One difference in accounting behavior between the OMCR model and 1:1 redundancy model is in the processing of the accounting session-time attribute which on the protecting node denotes the time when the host was instantiated on the protecting node.

In contrast, the session-time attribute in 1:1 redundancy model is recorded almost simultaneously on both BNG nodes at the time when the host is originally instantiated.

As a result, the session-time attribute is for the most part uninterrupted during the switchover in 1:1 model whereas in OMCR model, the session-time attribute is reset on the switchover to the protecting node.

10.2.11 Configuration Guidelines

- For all protected SRRP instances, the protected node should be the **preferred Master**. To achieve this, the SRRP priority should be higher in the protected node than in the protecting node. SRRP preemption is recommended in the protecting node to force it to become Master when possible.

Note: An SRRP switch from non-Master to Master in the protected node does not suffer the slow convergence observed when the non-Master to Master transition takes place in the protecting node. This is because the protected node always has the hosts instantiated in the data plane.

- ARP hosts configuration is strongly discouraged in protected group interfaces, unless the operator is ready to tolerate an incomplete redundancy mechanism for these hosts.
- SRRP tracking is strongly recommended to expedite the routing convergence upon an SRRP transition from non-Master to Master in the protecting node.
- It is recommended to have a 1:1 relationship between SRRP and subscriber subnets in order to have smooth routing advertisements based on SRRP state tracking.
- The use of M-SAPs should be preferred over the use of static SAPs. Static SAPs are supported in the OMCR mode but they are consuming resources in the protecting node even when the underlying SRRP instance is in a non-Master state.
- It is recommended that the capture-sap configuration include the **track-srrp** statement (at least for the protecting node). With this configuration the CPM will not process trigger packets when the leases cannot be created because the SRRP is not in the Master state. Configuring SRRP tracking at the capture-sap will offload the CPM from performing false authentication and MSAP creation attempts.
- Load balancing between Master and non-Master via export policies for SRRP must not be configured as the hosts are not instantiated in the protecting node when the corresponding SRRP state is non-Master.
- In order to minimize traffic impact in the event of node reboot, it is recommended to use **hold-time down ip** | **ipv6** *seconds* command under the subscriber-interface and allow enough time for the MCS database to reconcile. This is particularly important in the protected node. If the SRRP becomes master in the protected node before the database has been reconciled, the protecting node will remove the leases (non-Master state) which have not been synchronized. This would create partial outage.

- In order to avoid SRRP collisions, lack of resources and partial subscriber host instantiation, the use of fate-sharing-groups is not recommended. As long as an SRRP instance can be served by the protected node, it is preferred to keep it in the Master state in there, instead of switching it to the protecting node as part of the operation-group.

10.2.12 Troubleshooting Commands

Some of the commands that can assist in troubleshooting are listed below.

Note: To get a summary view of SRRPs and their OMCR status use the following command as shown below (the **domain** concept is reserved for future use):

```
*A:right-21# show redundancy multi-chassis omcr all
===============================================================================
Domain Table
===============================================================================
Domain    Domain SRRP  SRRP    Domain Instan. Failed Failed Reason
name       state ID    State    Color  Failed  Hosts
-------------------------------------------------------------------------------
N/A       N/A    201   Standby N/A    not-act 0
N/A       N/A    202   Standby N/A    not-act 0
N/A       N/A    203   Standby N/A    not-act 0
N/A       N/A    204   Standby N/A    not-act 0
N/A       N/A    301   Standby N/A    not-act 0
N/A       N/A    302   Standby N/A    not-act 0
N/A       N/A    303   Standby N/A    not-act 0
N/A       N/A    304   Standby N/A    not-act 0
N/A       N/A    401   Standby N/A    not-act 0
N/A       N/A    402   Standby N/A    not-act 0
N/A       N/A    403   Standby N/A    not-act 0
N/A       N/A    404   Standby N/A    not-act 0
N/A       N/A    501   Standby N/A    not-act 0
N/A       N/A    601   Standby N/A    not-act 0
N/A       N/A    701   Standby N/A    not-act 0
N/A       N/A    801   Standby N/A    not-act 0
N/A       N/A    901   Standby N/A    not-act 0
N/A       N/A    1001  Standby N/A    not-act 0
N/A       N/A    1101  Standby N/A    not-act 0
-------------------------------------------------------------------------------
No. of Entries: 19
===============================================================================
*A:right-21#
```

Note: To obtain specific SRRP OMCR information, OMCR information has been added to the **show srrp x detail** command:

```
*A:right-21# show srrp 1001 detail
===============================================================================
SRRP Instance 1001
===============================================================================
Description         : (Not Specified)
Admin State         : Up                 Oper State       : backupRouting
Oper Flags          : subnetMismatch
Preempt             : yes                One GARP per SAP : no
Monitor Oper Group  : None
System IP           : 10.20.1.1
Service ID          : IES 2
Group If            : grp.Dut-J.1        MAC Address      : 00:00:61:ac:ac:0a
Grp If Description  : N/A
Grp If Admin State  : Up                 Grp If Oper State: Up
Subscriber If       : ies-sub-if-svc-2
Sub If Admin State  : Up                 Sub If Oper State: Up
Address             : 10.1.0.1/16       Gateway IP       : 10.1.0.3
Address             : 10.2.0.1/16       Gateway IP       : 10.2.0.3
Msg Path SAP        : 8/2/2:2.4094
Admin Gateway MAC   : 00:00:51:ac:0a:01  Oper Gateway MAC : 00:00:51:ac:0a:01
Config Priority     : 1                  In-use Priority  : 1
Master Priority     : 100
Keep-alive Interval : 10 deci-seconds    Master Since     : 02/11/2014 11:38:52
Master Down Interval: 3.000 sec (Expires in 2.700 sec)
Fib Population Mode : all
VRRP Policy 1       : None               VRRP Policy 2    : None
OMCR Client status  : Sub-mgmt-ipoe
Instantiation failed: not-act            Failed IPOE Hosts: 0
OMCR Reason         :
```

Note: To have a view of the MCS synchronization including OMCR standby records:

```
*A:right-21# show redundancy multi-chassis sync peer 10.20.1.6 detail
===============================================================================
Multi-chassis Peer Table
===============================================================================
Peer
-------------------------------------------------------------------------------
Peer IP Address         : 10.20.1.6
Description             : (Not Specified)
Authentication          : Disabled
Source IP Address       : 10.20.1.1
Admin State             : Enabled
Warm standby            : Yes
Remote warm standby     : No
-------------------------------------------------------------------------------
Sync-status
-------------------------------------------------------------------------------
Client Applications     : SUBMGMT-IPOE SRRP
Sync Admin State        : Up
Sync Oper State         : Up
Sync Oper Flags         :
DB Sync State           : inSync
Num Entries             : 64026
```

```
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 64000
OMCR Alarm Entries      : 0
Rem Num Entries         : 64026
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
===============================================================================
MCS Application Stats
===============================================================================
Application             : igmp
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : igmpSnooping
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : subMgmtIpoe
Num Entries             : 64000
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 64000
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 64000
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : srrp
Num Entries             : 26
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 26
```

```
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : mcRing
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : mldSnooping
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : dhcpServer
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : subHostTrk
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : subMgmtPppoe
Num Entries             : 0
```

```
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : mcIpsec
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
Application             : mld
Num Entries             : 0
Lcl Deleted Entries     : 0
Alarm Entries           : 0
OMCR Standby Entries    : 0
OMCR Alarm Entries      : 0
-------------------------------------------------------------------------------
Rem Num Entries         : 0
Rem Lcl Deleted Entries : 0
Rem Alarm Entries       : 0
Rem OMCR Standby Entries: 0
Rem OMCR Alarm Entries  : 0
-------------------------------------------------------------------------------
===============================================================================
===============================================================================
Ports synced on peer 10.20.1.6
===============================================================================
Port/Encap                    Tag
-------------------------------------------------------------------------------
4/2/2
  2.1-2.4094                  Dut-F.1
===============================================================================
===============================================================================
DHCP Server instances synced on peer 10.20.1.6
===============================================================================
Router-Name                      Server-Name
  Tag
-------------------------------------------------------------------------------
No instances found
===============================================================================
===============================================================================
*A:right-21#
```

Note: To have the MCS database view of the sync status including OMCR status use the following command syntax:

```
*A:right-21# tools dump redundancy multi-chassis sync-database application sub-mgmt-
ipoe peer 10.20.1.6
The following totals are for:
 peer ip 10.20.1.6, port/lag ALL, sync-tag ALL, application SUBMGMT-IPOE
Valid Entries:                  64000
Locally Deleted Entries:        0
Locally Deleted Alarmed Entries: 0
Pending Global Delete Entries:  0
Omcr Alarmed Entries:           0
Omcr Standby Entries:           64000
*A:right-21#
```

10.3 OMCR Command Reference

10.3.1 Command Hierarchies

Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Basic System Configuration Guide* for information about redundancy CLI command descriptions and syntax.

config
- **redundancy**
  - **multi-chassis**
    - [**no**] **peer** *ip-address*
      - **warm-standby**

10.3.2 OMCR Configuration Commands

10.3.2.1 OMCR Configuration Commands

warm-standby

Syntax **warm-standby**

Context config>redundancy> multi-chassis>peer

Description This command enables Oversubscribed Multi-Chassis Redundancy (OMCR) model where subscriber hosts are synchronized between the two chassis only in the control plane and are kept there (as part of the Multi-Chassis Synchronization (MCS) state) until the switchover occurs. Link or nodal failure will trigger the switchover at which point the subscriber hosts are being fully instantiated in the control and the forwarding plane. This approach allows oversubscription of the resources in the central standby (or protecting) node that is backing-up a number of other active nodes. The total number of protected subscribers in the OMCR cluster exceeds the forwarding capacity of the protecting node. This is achievable by not fully occupying the resources for the subscriber hosts until the failure occurs.

The restoration times depend on the amount of the subscriber hosts that are affected by the switchover and it is related to the time needed for the full instantiation of the subscribers in the forwarding plane.

Although this command is configured on a peer level, the warm-standby property is a nodal characteristic. In other words, mixing of N:1 and 1:1 (hot standby) mode in the central standby node is not supported. Consequently all peers on the central standby node must be configured for warm-standby (N:1), or all peers must be configured for hot-standby (1:1) by omitting the warm-standby keyword from the configuration.

The peer of the central backup node is not aware of the redundancy model supported. In in other words, the peer of the central-backup node does not know whether it peers with a warm-standby peer or host-standby-peer. All nodes participating in this protection model must run SR OS R12.0 or higher.

11 WIFI Aggregation and Offload

11.1 WIFI Aggregation and Offload Overview

This solution set adds support for managing subscribers gaining network access over WLAN. The WLAN access enables a service provider to offer a mobile broadband service to its subscribers or to offload traffic on its or a partners macro cellular (3G/4G) network. The WLAN access can be from public hot-spots (indoor or outdoor APs), venues, enterprises, or home-spots (with public SSID).

The 7750 SR serves as a WLAN Gateway (WLAN-GW) providing Layer 3 termination and ESM for these subscribers. The connectivity from WLAN AP or AC can be over any existing access technology (DSL, PON, Fiber, DOCSIS, and so on), with Ethernet based connectivity from the access node (DSLAM, OLT, Eth MTU, Layer 2 CMTS) to the WLAN-GW. WLAN-GW functions could be on a standalone 7750 as shown in Figure 174 or could be an add-on functionality on existing 7750 based BNG as shown in Figure 175. WLAN connectivity to the WLAN-GW could be over a Layer 2 aggregation or an Layer 3 aggregation network (typical when WLAN-GW is upstream of an existing BNG or CMTS). In case of Layer 2 aggregation the connectivity to the WLAN-GW could be tagged or untagged Ethernet. In case of Layer 3 aggregation, supported connectivity option is Ethernet over GRE (or Eth-over-MPLS over GRE) tunnel originating from the AP/AC, and terminating on the WLAN-GW. The WLAN AP acts as a bridge, switching Ethernet frames into a GRE tunnel terminating on an MS-ISA in the WLAN-GW.

Figure 174 **Standalone WLAN-GW**

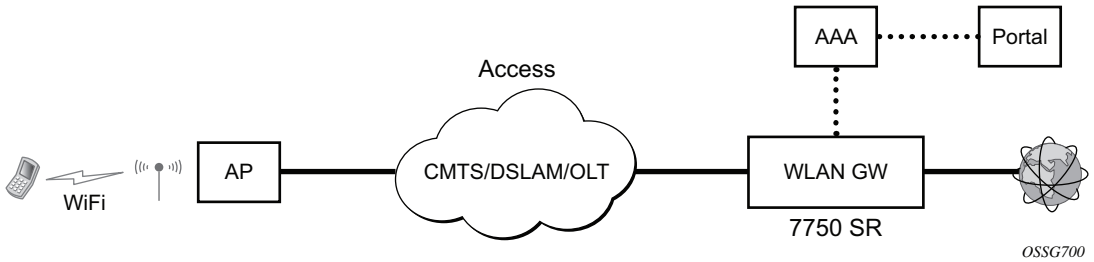

Figure 175 **WLAN-GW Functions on Existing BNG**

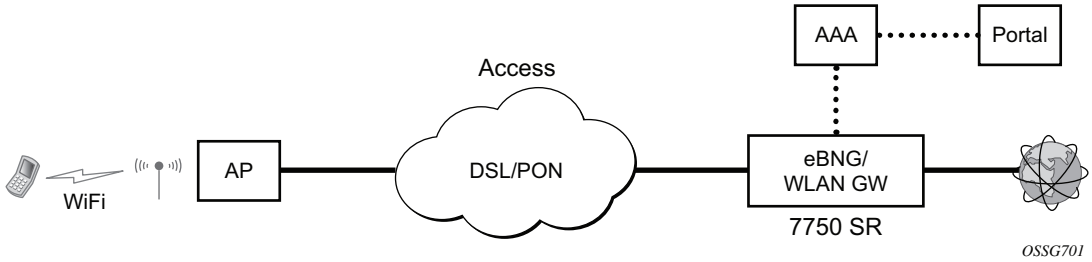

AP Connectivity to the WLAN-GW could be direct Ethernet (tagged or untagged) or could be Ethernet over GRE. With the bridged AP using GRE tunnels, the WLAN-GW solution elements are discussed in the following sections.

11.2 Layer 2 over Soft-GRE Tunnels

Soft-GRE refers to stateless GRE tunneling, whereby the AP forwards GRE encapsulated traffic to the WLAN-GW, and the GW reflects back the encapsulation in the downstream traffic towards the AP. WLAN-GW does not require any per-AP end-point IP address configuration. The WLAN-GW learns the encapsulation as part of creating the subscriber state on processing the encapsulated control and data traffic. Following are some of the advantages of soft-GRE:

- Resources are only consumed on the WLAN-GW if there is one or more active subscriber on the AP. Merely broadcasting an SSID from an AP does not result in any state on the WLAN-GW.
- No per-AP tunnel end-point configuration on WLAN-GW. This is important as the AP can get renumbered.
- No control protocol to setup and maintain tunnel state on WLAN-GW.

Soft-GRE tunnel termination is performed on dedicated IOMs with MS-ISAs (referred to as WLAN-GW IOM) Each slot requires two MS-ISAs dedicated for soft-GRE tunnel termination. MS-ISA provides tunnel encapsulation/decapsulation, bandwidth shaping per tunnel (or per-tunnel per SSID), and anchor point for inter-AP mobility. The ESM function such as per-subscriber anti-spoofing (IP and MAC), filters, hierarchical policing, and lawful intercept are provided on the carrier IOM corresponding to the ISA where the subscriber is anchored.

11.2.1 Encapsulation

The GRE encapsulation is based on RFC 1701/2784, *Generic Routing Encapsulation (GRE)*, WLAN-GW encapsulates according to RFC 1701 with all the flag fields set to 0, and no optional fields present. WLAN-GW is able to receive both encapsulation specified in RFC 1701 and RFC 2784, with all flag fields set to 0, and no optional fields present in the header.

Figure 176 **Encapsulation Example**

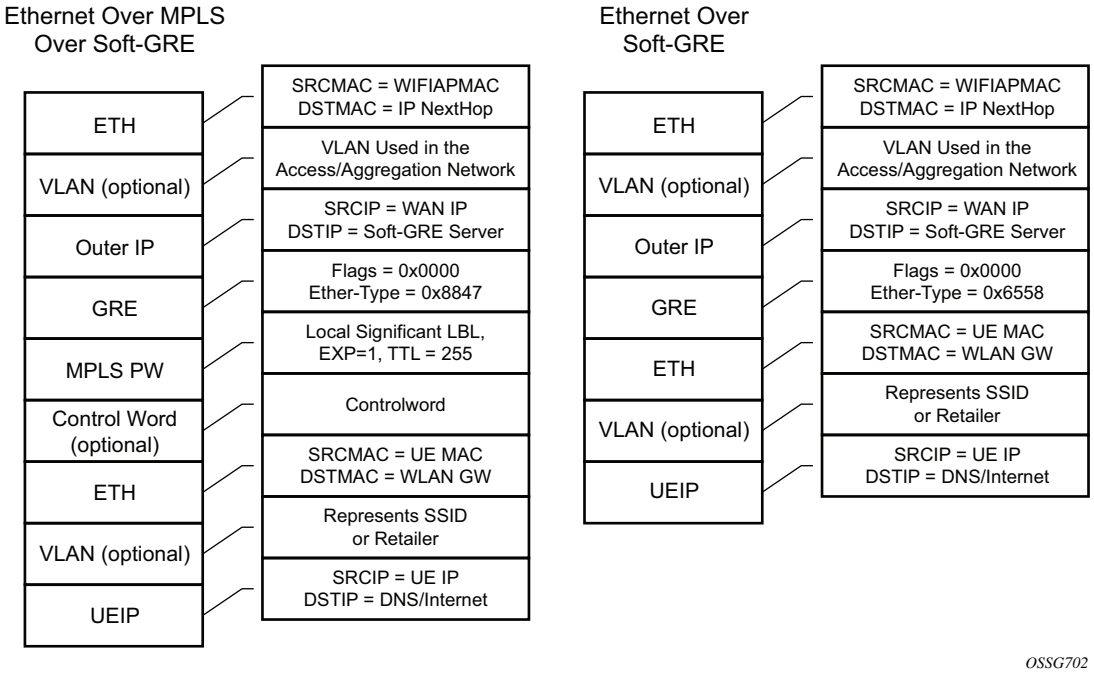

The encapsulation is built as follows:

- Outer Ethernet header: (14 bytes)
  - Source MAC: MAC address of the WIFI AP/RG/HGW HW address
  - Destination MAC: MAC address of the first IP NH the WIFI AP/RG/HGW is connected to (for example, CMTS, IP aggregation router, BNG, and so on)
- Outer VLAN: (4 bytes): optional, typically used for service delineation in the access or aggregation network.
- Outer IPv4 Header: (20 bytes)
  - Source IP — IP address used for WAN addressing which is retrieved by the AP/RG from the ISP through DHCP, PPPoX, and so on

  - Destination IP — Soft-GRE server address which can be retrieved by a DHCP Option, PPPoX option or configured by TR69 or configured statically in a boot file (in cable environment).
  - DSCP — Reflects QoS used in the access/aggregation network.
  - TTL — Should be set to 255 or should reflect the amount of IP hops in the access/aggregation network
- GRE: (4 bytes)
  - All flags are set to 0, such as checksum, sequence number and keys are not present.
  - The Ether-Type is set to 0x6558 for native Ethernet is used, and 0x8847 when MPLS encapsulation is used.
- MPLS Pseudowire Label (4 bytes)
  - Label Value, statically assigned in the WiFi AP/Controller and reflected back from the soft-GRE server to the WIFI AP/Controller. The Label is unique within the context of the source IP address of the tunnel.
  - EXP: 0 (not used)
  - TTL: 255 (not used)
- Inner Ethernet header: (14 bytes)
  - Source MAC: MAC address of the UE
  - Destination MAC: MAC address of the soft-GRE server/WLAN-GW.
- Inner VLAN: (4 bytes): optional, inserted by AP/RG per unique SSID (typically, when the AP is providing SSID per retailer). WLAN-GW allows mapping the VLAN to a service context per retailer, in the data plane.
- Inner IPv4 Header: (20 bytes)
  - Source IP: Client’s IP address obtained via DHCP (tunneled).
  - Destination IP: IP address of the destination client trying to reach.
  - DSCP: set by the client/application
  - TTL: set by the client/application

Soft-GRE tunnel termination is performed on dedicated IOMs with MS-ISAs (referred to as WLAN-GW IOM). Each WLAN-GW IOM requires both MS-ISAs to be plugged in for soft-GRE tunnel termination. MS-ISA provides tunnel encapsulation/ decapsulation and anchor point for inter-AP mobility. The carrier IOMs of the ISA where the tunnel is terminated performs bandwidth shaping per tunnel (or per-tunnel per SSID). ESM function such as per-subscriber anti-spoofing (IP and MAC), filters, hierarchical policing, and lawful intercept are provided on the carrier IOM corresponding to the ISA where the subscriber is anchored.

N:M warm standby redundancy is supported for WLAN-GW IOM slots. One WLAN-GW group can be configured with set of WLAN-GW IOMs, and a limit of active IOMs. Incoming soft-GRE tunnel contexts and corresponding subscribers are load-balanced amongst the MS-ISAs on active IOMs. Tunnel load-balancing is based on outer source IP address of the tunnel. Subscriber load-balancing is based on UE's MAC address in the source MAC of the Ethernet payload in the tunnel. IOM(s) beyond the active limit act as warm standby, and take over the tunnel termination and subscriber management functions from failed WLAN-GW slot.MS-ISAs on WLAN-GW IOMs can also be configured to perform NAT function.

```
config isa wlan-gw-group <group-id>
   [no] active-iom-limit <number>
   [no] description <description-string>
   [no] distributed-sub-mgmt
        [no] isa-aa-group <aa-group-id>
   [no] * iom <slot-number>
      nat
        [no] radius-accounting-policy <nat-accounting-policy>
        [no] session-limits
          [no] reserved <num-sessions>
          [no] watermarks high <percentage> low <percentage>
   [no] shutdown
```

An ESM and soft-gre configuration is required for wlan-gw functions. Subscriber and group interfaces are configured as part of normal ESM configuration. The group interface is enabled for wlan-gw by configuration. L2oGRE is the currently supported soft tunnel types. The wlan-gw related configuration includes the following:

- Tunnel end-point IP address.
- Service context for tunnel termination.
- TCP MSS segment size. This is set in TCP SYN and SYN-ACKs by wlan-gw to adjust to the MTU on access/aggregation network in order to prevent fragmentation of upstream and downstream TCP packets.
- Mobility related configuration, including mobility trigger packet types (normal data or special Ethernet IAPP fame), and hold-down time between successive mobility triggers.
- VLAN to retailer mapping. The AP typically inserts a unique dot1Q tag per retail service provider in the Ethernet payload. The mapping of dot1Q tag to retail service context is configured under wlan-gw tunnel. The subscriber is then created in the configured retail service context. The retail service context can also be provided by AAA server in authentication-accept message based on subscriber credentials or SSID information contained in DHCP Option82.

- Egress QoS configuration for downstream traffic entering the wlan-gw module for tunnel encapsulation. This includes type of aggregate bandwidth shaping (per-tunnel or per-retailer), aggregate-rate-limit, egress QoS policy and scheduler policy. The tunnel shaping can be configured to be applied only when there is more than one subscriber on the tunnel. By default the shaping if configured is applied when first subscriber on the tunnel logs in.

```
*B:Dut-C>config>service>vprn>sub-if>grp-if>wlan-gw# info detail
----------------------------------------------
                        authentication
                            no authentication-policy
                            hold-time sec 5
                        exit
                        no data-triggered-ue-creation
                        dhcp
                            shutdown
                            active-lease-time min 10
                            initial-lease-time min 10
                            no l2-aware-ip-address
                            no primary-dns
                            no primary-nbns
                            no secondary-dns
                            no secondary-nbns
                        exit
                        egress
                            no agg-rate-limit
                            no hold-time
                            qos 1
                            no scheduler-policy
                            no shape-multi-client-only
                            no shaping
                        exit
                        gw-address 10.1.1.4
                        no gw-ipv6-address
                        no http-redirect-policy
                        no nat-policy
                        mobility
                            hold-time 5
                            no trigger
                        exit
                        router 70
                        no tcp-mss-adjust
                        track-mobility
                            mac-format "aa:"
                            no radius-proxy-cache
                        exit
                        wlan-gw-group 3
                        vlan-tag-ranges
                            range start 0 end 100
                                authentication
                                    no authentication-policy
                                    hold-time sec 5
                                exit
                                no data-triggered-ue-creation
                                dhcp
                                    shutdown
                                    active-lease-time min 10
```

```
                initial-lease-time min 10
                no l2-aware-ip-address
                no primary-dns
                no primary-nbns
                no secondary-dns
                no secondary-nbns
            exit
            no http-redirect-policy
            no nat-policy
            retail-svc-id 35
            track-mobility
                mac-format "aa:"
                no radius-proxy-cache
            exit
        exit
    exit
    no shutdown
```

11.2.2 Data Path

In the upstream direction, the ingress IOM receiving the GRE tunneled packets from the WIFI AP or AC, load-balances tunnel processing amongst the set of MS-ISAs on the active WLAN-GW IOMs in the WLAN-GW group. The load-balancing is based on a hash of source IP address in the outer IP header. The MS-ISA receiving the GRE encapsulated packets removes the tunnel encapsulation, and internally tunnels (MAC-in-MAC, using BVPLS) the packet to an anchor MS-ISA on the WLAN-GW IOM. All traffic from a given UE is always forwarded to the same anchor MS-ISA based on hashing on UE's MAC address. The MS-ISA provides a mobility anchor point for the UE. The UE MAC's association to the GRE tunnel identifier is created or updated. The corresponding IOM provides ESM functions including ESM lookup, ingress ACLs and QoS. DHCP packets are forwarded to the CPM from the anchor IOM.

In the downstream direction, the IP packets are forwarded as normal from the network IOM (based on route lookup yielding subscriber subnet) to the IOM where the ESM host is anchored. ESM processing including per UE hierarchical policing and LI is performed on the anchor IOM. Configured MTU on the group-interface is enforced on the IOM, and if required packets are fragmented. The packets are then forwarded to the appropriate anchor MS-ISA housed by this IOM. Lookup based on UE's MAC address is performed to get the tunnel identification, and the packets are MAC-in-MAC tunneled to the MS-ISA terminating the GRE tunnel. Aggregate shaping on the tunneled traffic (per tunnel or per retailer) is performed on the carrier IOM housing the tunnel termination MS-ISA. The tunnel termination MS-ISA removes MAC-in-MAC encapsulation, and GRE encapsulates the Layer 2 packet, which exits on the Layer 3 SAP to the carrier IOM. The GRE tunneled packet is forwarded to the right access IOM toward the WIFI AP based on a routing lookup on IP DA in the outer header.

11.2.3 MDA-based Redundancy

The WLAN-GW supports N:M warm standby redundancy on an MDA level. With this feature, it is possible to configure separate MDA's in the WLAN-GW group without the restriction that all MDA's in an IOM be BB modules. Up to 14 BB modules can be configured and active at one time.

MDA-based redundancy makes hardware provisioning more flexible, but does not guarantee that all ESM session states is recovered after an MDA failure. Because BB MDA's share IOM resources with another MDA in the same IOM, there is no guarantee that the same set of resources is available on the IOM of the new MDA after a failure. This restriction does not apply to DSM.

There are several ways to ensure that the system will have adequate IOM resources following a failure.

- Use symmetric provisioning over all IOMs with BB modules. For example, do not create a 2:3 active:standby group by putting two BB MDA's on one IOM (IOM A) and one BB MDA on another IOM (IOM B). When the active ISAs are spread over IOM A and IOM B, the WLAN-GW has the full resources of two IOMs. However, after a failure, it is possible that both active ISAs is on IOM A and limited to the resources of a single IOM.
- Combine BB MDAs with MDAs and/or services using predictable and fixed resources. For example, combining BB functionality with AA modules will provide enhanced traffic services because AA modules typically have a fixed resource usage that allows you to predict resources following a failure.
- Create a safety buffer by leaving some resources unallocated. By not maximizing resource usage on a single IOM, it is more likely those resources will also available on a backup location.

For more details, contact a Nokia representative.

11.3 Tunnel Level Egress QoS

Downstream traffic can be subjected to aggregate rate-limit per tunnel or per tunnel and per retailer combination (in case of wholesale). Typically a unique SSID is used per retailer for wholesale on the AP, and is reflected via unique dot1Q tag. With a wlan-gw tunnel per AP, the tunnel encapsulation is performed on the tunnel ISA. The downstream traffic on the tunnel IOM is received over B-VPLS from the anchor IOM, and is MAC-in-MAC (802.1ah) encapsulated. I-SID in the packet represents the GRE tunnel or tunnel and retailer combination. SAP-egress QoS policy defining queues (with rates), and FC to queue mapping, can be specified under the wlan-gw interface. This policy is applicable to all tunnels (or tunnel and SSID combinations) associated with the wlan-gw interface, and is attached to corresponding I-SIDs on the B-VPLS SAP. Traffic is shaped into these queues based on configured queue rates. An aggregate rate-limit applied across queues on an I-SID (representing tunnel or tunnel and retailer combination) can be configured under the wlan-gw interface (represented by the wlan-gw node under the group-interface configuration). The aggregate rate-limit works in conjunction with a port-scheduler. The port-scheduler corresponds to the internal port between tunnel ISA and its carrier IOM, and is specified at the wlan-gw IOM group level. The rate-limit includes the B-VPLS encapsulation overhead. The configuration is shown in Figure 178. Queues per I-SID also work with virtual-scheduler (with or without a port scheduler). Virtual-scheduling and aggregate-rate enforcement are mutually exclusive. Configuration is shown in Figure 179. Egress SAP QoS policy, aggregate rate-limit, port-scheduler, and virtual-schedulers are described in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Quality of Service Guide*. The SAP egress QoS policy associated with a wlan-gw interface implicitly creates queues (and scheduler association) on ISIDs as corresponding wlan-gw tunnels are created. General ISID queuing and shaping is defined in the SR OS Services Guide.

A configuration node under wlan-gw interface (egress) controls where the egress shaping is applied, and can specify either tunnel or retailer (tunnel and retailer combination in case of wholesale). Per I-SID shaping resources can be held after the last subscriber on the tunnel is deleted, for a configurable amount of time (hold-time) configured under the wlan-gw interface. During ISA or IOM failover the tunnel resources on the IOM kept due to hold-time are reclaimed. ISID shaping can be configured (via knob shape-multi-client) to be applied only when there is more than one UE on the corresponding tunnel (or tunnel and retailer combination). A total of 40,000 shaped tunnels (or shaped tunnel & retailer combinations) are supported per WLAN-GW IOM. Hardware resources for tunnel (ISID) shapers are shared with subscribers. With 3 WLAN-GW IOMs per chassis, a maximum of 98,000 (3 *64K / 2) shaped tunnels and subscribers can be supported per chassis.

The following output depicts per tunnel or per tunnel/SSID egress QoS (with aggregate-rate and port-scheduler).

// Port-scheduler

```
config>qos#
      port-scheduler-policy "lo-gre-port-sched"
           max-rate 5000
           level 1 rate 1000 cir-rate 1000
           level 8 rate 500 cir-rate 500
      exit
exit
```

// Egress queues (per ISID) parented by port-scheduler specified under associated wlan-gw interface

```
config>qos>
   sap-egress 3 create
       queue 1 create
          rate 300
          port-parent level 1 weight 10 cir-level 1 weight 10
       exit
       queue 2 create
          rate 100
          port-parent level 8 weight 10 cir-level 8 weight 10
       fc af create
           dot1p 2
           de-markweight
       exit
       fc be create
          queue 1
          dot1p 0
          de-mark
       exit
       fc ef create
           queue 2
           dot1p 5
           de-mark
       exit
   exit
exit
```

// The wlan-gw interface refers to SAP egress QoS policy and aggregate rate-limit for associated ISIDs

```
config>service>ies>sub-if>grp-if>wlan-gw>egress
      agg-rate-limit 2000
      hold-time 300
      qos 3
      shaping per-tunnel
      shape-multi-client
exit
```

// Port-scheduler parenting queues (per ISID)

Per Tunnel or Per Tunnel/SSID Egress QoS (with aggregate-rate and port-scheduler)

```
config>isa>wlan-gw-group#
      active-iom-limit 1
      tunnel-port-policy " lo-gre-port-sched "
      iom 2
      iom 3
      no shutdown
exit
```

The following output depicts per tunnel or per tunnel/SSID egress QoS (with virtual-scheduler).

```
// hierarchical virtual scheduler
config>qos#
      scheduler-policy “virtual-sched-policy”
           tier1
               scheduler “all-traffic” create
                    rate 10000
               exit
           exit
           tier2
               scheduler “non-voice” create
                   parent all-traffic cir-level 1
                   rate 9000
               exit
               scheduler “voice” create
                  parent all-traffic level 2 cir-level 2
                  rate 3000
               exit
           exit

       exit
```

// egress queues (per ISID) parented by virtual scheduler

```
config>qos>
   sap-egress 3 create
       queue 1 create
          parent “non-voice”
          rate 2000 cir 1000
       exit
       queue 2 create
          parent “voice”
          rate 500 cir-rate 500
       fc be create
          queue 1
          dot1p 0
          de-mark
       exit
       fc ef create
           queue 2
           dot1p 5
           de-mark
       exit
   exit
exit
```

// A wlan-gw interface refers to SAP egress QoS policy and hierarchical scheduler for associated ISIDs

```
config>service>ies>sub-if>grp-if>wlan-gw>egress
      hold-time 300
      qos 3
      scheduler-policy "virt-sched-policy"
      shaping per-tunnel
      shape-multi-client
exit
```

11.3.1 QoS Overrides

Per-tunnel QoS overrides can be provided in order to have more specific values for certain tunnels. This allows the WLAN gateway to accommodate a heterogeneous set of access points in one WiFi network; for example, provisioning different bandwidth for homespots and hotspots. For a detailed list of allowed overrides and values, refer to the *7750 SR and VSR RADIUS Attributes Reference Guide*.

QoS overrides can be signaled via RADIUS in UE authentication and in UE CoA. The override is applied to the tunnel where the UE is currently active. Therefore, two scenarios should be considered:

- If the first UE on a tunnel is also a new UE that is authenticated, any tunnel QoS override can be signaled in the Access-Accept message for that UE.
- If the first UE on a tunnel is a previously authenticated UE that moved from another tunnel, a mobility triggered Interim Update can be sent and used as a trigger to send a CoA with the per-tunnel QoS overrides.

Figure 177 shows flows for both scenarios.

Figure 177 **QoS Override Flows**

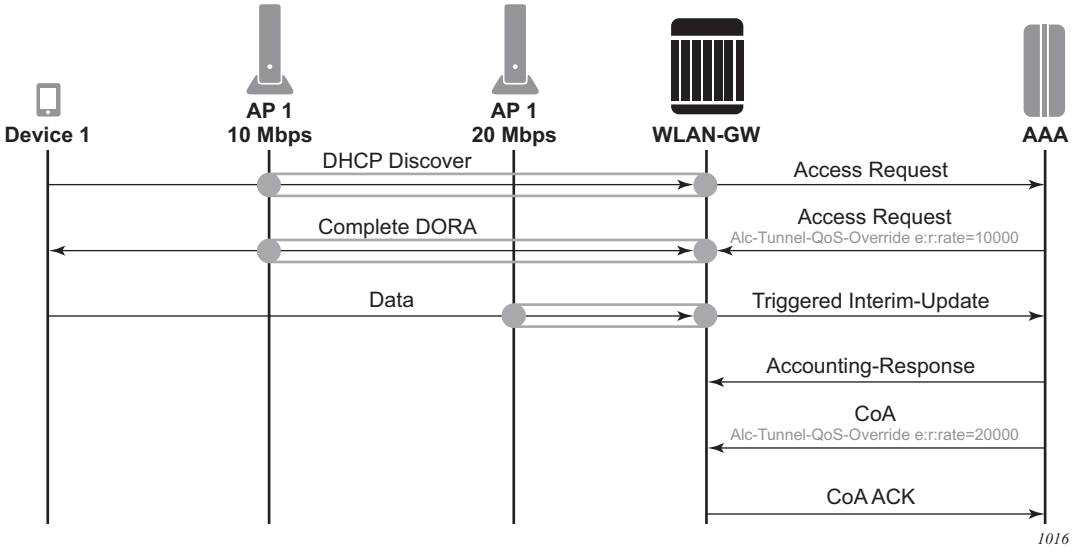

To guarantee that the CoA is applied to the correct tunnel, use the CoA key consisting of the *nas-port-id* and the UE IP address; both can be derived from the interim update. In the WLAN gateway, the *nas-port-id* identifies the tunnel to which the UE is linked, and, if a mismatch is detected, the CoA is not applied. Other CoA keys are also supported but do not offer this guarantee.

11.3.2 Operational Commands

Egress per tunnel (or per tunnel, per SSID) QoS with aggregate rate-limit and port-scheduler.

```
show router 50 wlan-gw soft-gre-tunnels detail
===============================================================================
Soft GRE tunnels
===============================================================================
Remote IP address           : 239.1.1.2
Local IP address            : 10.1.1.1
ISA group ID                : 1
ISA group member ID         : 1
Time established            : 2012/06/19 20:31:36
Number of UE                : 1

Tunnel QoS
----------
Operational state           : active
Number of UE                : 1
Remaining hold time (s)     : N/A
Service Access Points(SAP)
===============================================================================
```

```
Service Id         : 2147483650
SAP                : 2/1/lo-gre:1             Encap             : q-tag
Description        : Internal SAP
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 06/19/2012 07:13:31
Last Mgmt Change   : 06/19/2012 20:30:24
-------------------------------------------------------------------------------
Encap Group Specifics
-------------------------------------------------------------------------------
Encap Group Name   : _tmnx_SHAPER_GR000       Group Type        : ISID
Qos-per-member     : TRUE
Members            :1
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
E. qos-policy      : 3                        Q Frame-Based Acct: Disabled
E. Sched Policy    :                          E. Agg-limit      : 4000
-------------------------------------------------------------------------------
Encap Group Member 1 Base Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

Forwarding Engine Stats
                        Packets                 Octets
For. InProf           : 0                       0
For. OutProf          : 0                       0
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
-------------------------------------------------------------------------------
Encap Group Member 1 Queue Statistics
-------------------------------------------------------------------------------
                        Packets                 Octets
Egress Queue 1
For. InProf           : 0                       0
For. OutProf          : 0                       0
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
===============================================================================
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================


show qos scheduler-hierarchy sap 2/1/lo-gre:1 encap-
group "_tmnx_SHAPER_GR000" member 1 detail
===============================================================================
Scheduler Hierarchy - Sap 2/1/lo-gre:1
===============================================================================
Egress Scheduler Policy :
-------------------------------------------------------------------------------
Legend :
(*) real-time dynamic value
(w) Wire rates
B   Bytes
-------------------------------------------------------------------------------
Root (Egr)
| slot(2)
```

```
|--(Q) : -2147483646->2/1/lo-gre:1->EG(_tmnx_SHAPER_GR000):1->1  (Port 2/1/lo-
gre Orphan)
|   |    AdminPIR:10000000   AdminCIR:0
|   |    AvgFrmOv:100.00
|   |    AdminPIR:10000000(w) AdminCIR:0(w)
|   |    CBS:0 B             MBS:12582912 B
|   |    Depth:0 B           HiPrio:1376256 B
|   |    MaxAggRate:4000(w)    CurAggRate:0(w)
|   |
|   |    [Within CIR Level 0 Weight 0]
|   |    Assigned:0(w)       Offered:0(w)
|   |    Consumed:0(w)
|   |
|   |    [Above CIR Level 1 Weight 0]
|   |    Assigned:4000(w)    Offered:0(w)
|   |    Consumed:0(w)
|   |
|   |    TotalConsumed:0
|   |    OperPIR:4000        OperCIR:0
|   |
|   |    PktByteOffset:add 0*
|   |    OnTheWireRates:false
|   |    ATMOnTheWireRates:false
|   |    LastMileOnTheWireRates:false
```

Egress per tunnel (or per tunnel, per SSID) QoS with hierarchical virtual scheduler.

```
show router 50 wlan-gw soft-gre-tunnels detail
===============================================================================
Soft GRE tunnels
===============================================================================
Remote IP address           : 239.1.1.2
Local IP address            : 10.1.1.1
ISA group ID                : 1
ISA group member ID         : 1
Time established            : 2012/06/19 20:43:03
Number of UE                : 1

Tunnel QoS
----------
Operational state           : active
Number of UE                : 1
Remaining hold time (s)     : N/A
Service Access Points(SAP)
===============================================================================
Service Id         : 2147483650
SAP                : 2/1/lo-gre:1             Encap             : q-tag
Description        : Internal SAP
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 06/19/2012 07:13:31
Last Mgmt Change   : 06/19/2012 20:30:24
-------------------------------------------------------------------------------
Encap Group Specifics
-------------------------------------------------------------------------------
Encap Group Name   : _tmnx_SHAPER_GR000       Group Type        : ISID
```

```
Qos-per-member     : TRUE
Members            :
1
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
E. qos-policy      : 3                        Q Frame-Based Acct: Disabled
E. Sched Policy    : virtual_scheduler_policy E. Agg-limit      : -1
-------------------------------------------------------------------------------
Encap Group Member 1 Base Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

Forwarding Engine Stats
                        Packets                 Octets

For. InProf           : 2                       752
For. OutProf          : 0                       0
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
-------------------------------------------------------------------------------
Encap Group Member 1 Queue Statistics
-------------------------------------------------------------------------------
                        Packets                 Octets

Egress Queue 1
For. InProf           : 2                       752
For. OutProf          : 0                       0
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
===============================================================================
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================


show qos scheduler-hierarchy sap 2/1/lo-gre:1 encap-
group "_tmnx_SHAPER_GR000" member 1 detail
===============================================================================
Scheduler Hierarchy - Sap 2/1/lo-gre:1
===============================================================================
Egress Scheduler Policy :
-------------------------------------------------------------------------------
Legend :
(*) real-time dynamic value
(w) Wire rates
B   Bytes
-------------------------------------------------------------------------------
Root (Egr)
| slot(2)
|--(S) : virtual_scheduler (Port 2/1/lo-gre)
|   |    AdminPIR:4000       AdminCIR:0(sum)
|   |
|   |    AvgFrmOv:105.31(*)
|   |    AdminPIR:4212(w)    AdminCIR:0(w)
|   |
|   |    [Within CIR Level 0 Weight 0]
|   |    Assigned:0(w)       Offered:0(w)
|   |    Consumed:0(w)
```

```
|   |
|   |    [Above CIR Level 1 Weight 1]
|   |    Assigned:4212(w)    Offered:0(w)
|   |    Consumed:0(w)
|   |
|   |
|   |    TotalConsumed:0(w)
|   |    OperPIR:3999
|   |
|   |    [As Parent]
|   |    Rate:3999
|   |    ConsumedByChildren:0
|   |
|   |
|   |--(Q) : -2147483646->2/1/lo-gre:1->EG(_tmnx_SHAPER_GR000):1->1
|   |   |    AdminPIR:10000000   AdminCIR:0
|   |   |    AvgFrmOv:105.31(*)
|   |   |    CBS:0 B             MBS:12582912 B
|   |   |    Depth:0 B           HiPrio:1376256 B
|   |   |
|   |   |    [Within CIR Level 0 Weight 1]
|   |   |    Assigned:0          Offered:0
|   |   |    Consumed:0
|   |   |
|   |   |    [Above CIR Level 1 Weight 1]
|   |   |    Assigned:3999       Offered:0
|   |   |    Consumed:0
|   |   |
|   |   |    TotalConsumed:0
|   |   |    OperPIR:4000        OperCIR:0
|   |   |
|   |   |    PktByteOffset:add 0*
|   |   |    OnTheWireRates:false
|   |   |    ATMOnTheWireRates:false
|   |   |    LastMileOnTheWireRates:false
```

11.4 Authentication

The solution supports multiple authentication mechanisms. Type of authentication support depends on the WIFI AP, UE capabilities and customer preference. In case of 802.1x/EAP capable WIFI APs, supporting secure SSIDs via 802.11i/WPA2, various EAP based authentication such as SIM/uSIM based (SIM/AKA/AKA'), TTLS, PEAP, certs, and so on, are supported. The solution also supports web-portal based authentication with or without WISPr client on the UE. EAP and portal authentication works independent of the type of connectivity from the AP (tunneled or native IP).

The SR OS WLAN-GW uses the IPoE session concept to authenticate and manage UEs in ESM. Every WLAN-GW group interface uses a pre-defined default **ipoe-session** policy that cannot be changed or disabled. The contents of the default policy also cannot be changed and always uses **sap-mac** as a key. The **ipoe-session session-timeout** can optionally be ignored in a **wlan-gw** context. This is to support closed SSID authentication where the session-timeout is relative to the last re-authentication while for **ipoe-session** the timeout is absolute to the start of the session.

It can be useful to identify the AP and the SSID to which a UE is connected. Therefore, the AP MAC and SSID name can be learned as follows:

- From the called-station-id as defined in RFC 3580, *Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.1 Certificate Handling*
- From DHCP circuit-Id or DHCPv6 interface-ID, if those options use the format specified below
- From ARP or ND over GRE as specified in section 11.10. This only identifies the AP MAC, not the SSID
- From the L2TPv3 cookie as specified in section 11.22. This only identifies the AP MAC, not the SSID

The format used for DHCP(v6) is AP-MAC;SSID-STRING;SSID-Type, where the AP-MAC should contain the AP MAC address in colon separated format (xx:xx:xx:xx:xx:xx), the SSID string should not contain the ";" delimiter and the SSID type is a single character "s" (secure) or "o" (open).

For example, if AP-MAC is "00:10:A4:23:19:C0", SSID is "SP1-wifi", and SSID-type is secure, then the value of circuit-id or interface-id would be the string "00:10:A4:23:19:C0;SP1-wifi;s".

11.4.1 EAP-Based Authentication

In this model the WIFI AP supports a RADIUS client, and originates RADIUS messages based on 802.1x/EAP exchange with the UE. It sends EAP payload in RADIUS messages towards the RADIUS server or RADIUS proxy. 7750 WLAN-GW can be configured as a RADIUS proxy for the WIFI APs. The WIFI AP should be configured with the IP address of the RADIUS proxy, and should send authentication and accounting messages non-tunneled, natively routed to the RADIUS proxy. See Figure 178.

The RADIUS proxy function allows 7750 SR to look at the RADIUS authentication and accounting messages and create or update corresponding subscriber state. RADIUS proxy transparently forwards RADIUS messages between AP (authenticator) and the AAA server. The access-request message contains standard RADIUS attributes (including user-name), and the EAP payload. Standard authentication algorithms negotiated with EAP involve multiple round-trips (challenge/response) between AP (and UE) and the AAA server.

Once authentication is complete, AAA server passes back subscriber related configuration parameters as well as the computed session keys (aka pair-wise master key) for 802.11i to the AP. These keys are encrypted using shared secret between AP (authenticator) and the AAA server. 7750 WLAN-GW can optionally cache authentication information of the subscriber from access-request and access-accept messages. The cached information allows local authorization of subsequent DHCP messages from the UEs behind the AP against the cached state on the 7750 RADIUS proxy, and avoids another trip to the RADIUS server.

Figure 178 **EAP Authentication Call Flow with WLAN-GW RADIUS Proxy**

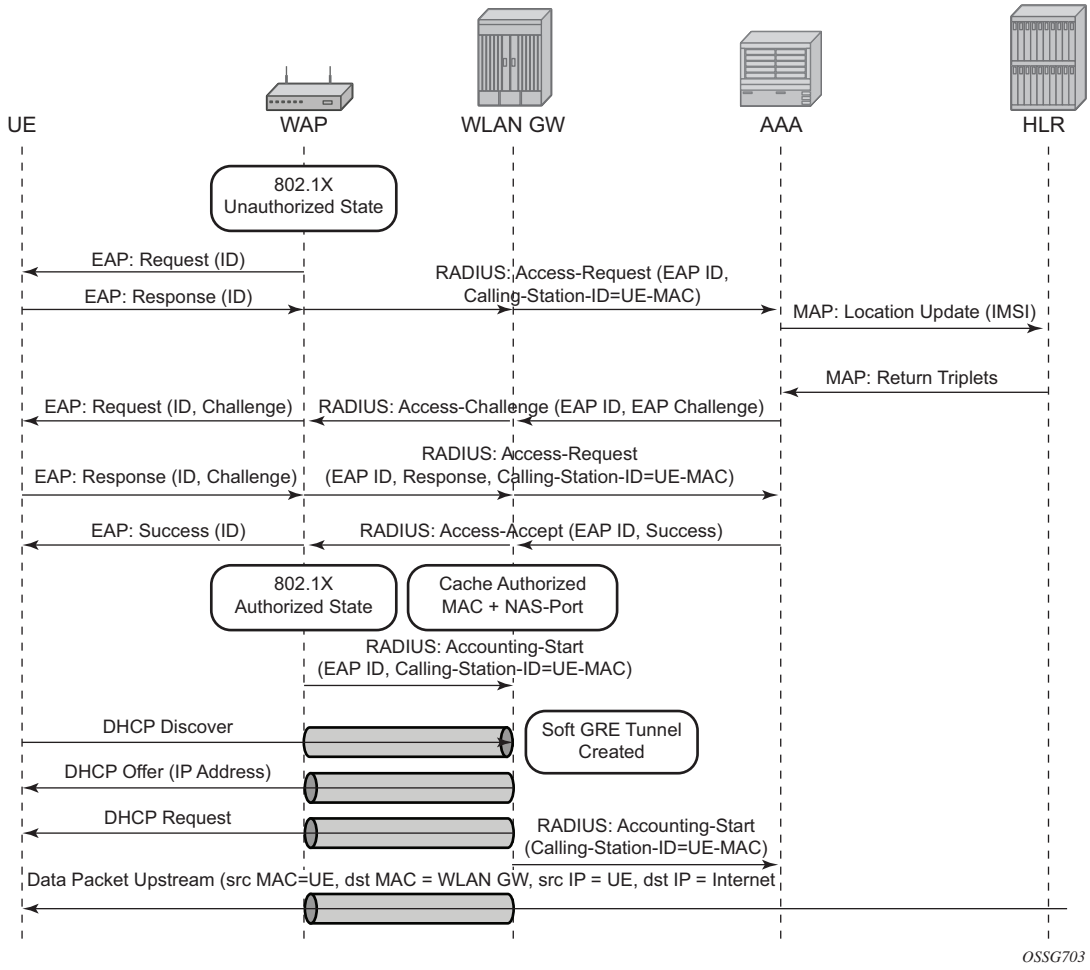

11.4.1.1 RADIUS Proxy

RADIUS proxy can be configured per service router (base or VPRN). The proxy acts as a server towards the WIFI AP RADIUS clients, and as a client towards RADIUS server(s). Therefore, both client and server parts of the RADIUS proxy need to be configured. The attribute from access-request or response message that serves as the key for the cache is configurable. The key configuration is mandatory for enabling the cache. Commonly the key is the MAC address of the UE, which is available in subsequent DHCP request, and used to locate the cache entry. The UE's MAC address is typically available in the Calling-station-Id attribute (31) in the RADIUS access-request message from the AP. The proxy can be configured for both authentication and accounting. The radius server policies referred by RADIUS proxy are configured under "aaa" context. If caching is enabled in the RADIUS proxy, the

subscriber attributes returned in access-accept are cached. These can include 802.1x credentials/keys, IP address or pool, DNS information, default gateway information, retail-service-id, SLA-profile, filter parameters, charging information, session keys (MS-MPPE-RECV-KEY, MS-MPPE-SEND-KEY) and so on. If subsequent DHCP DISCOVER is not received within the configured timeout, the cache entry is removed.

The following output displays a RADIUS proxy configuration.

```
config>service>ies>
config>service>vprn>
     description "Default Description For VPRN ID 50"
        interface "listening_radius_server" create
          address 10.9.9.9/32
          loopback
        exit

     radius-proxy
         server "radius_proxy" purpose accounting authentication create
             cache
                 key packet-type request attribute-type 31
                 timeout min 5
                 track-accounting stop interim-update accounting-on accounting-off
                 no shutdown
             exit
             default-accounting-server-policy "radius_acct_server_policy"
             default-authentication-server-policy "radius_Auth_server_policy"
             interface "listening_radius_server"
             load-balance-key attribute-type 102 vendor 5
             secret "AQepKzndDzjRI5g38L3LbbN3E8qualtn" hash2
             send-accounting-response
             no shutdown
         exit
```

11.4.1.1.1 RADIUS Proxy — Server Load-Balancing

RADIUS proxy can be configured for load-balancing to multiple authentication and accounting servers. Load-balancing can be round-robin or hash based, and is configured via access-algorithm under RADIUS policy. With round-robin the first RADIUS request is sent to the first server, the second request to the second server and so on. With hash, it is possible to load-balance subscribers across a set of servers. Based on the configured hash key, configured in the RADIUS proxy, it can be ensured that all RADIUS messages for a single subscriber are sent to the same server. The hash key can include any specified standard or vendor-specific RADIUS attribute. An example is calling-station-id which contains subscriber's MAC address).

If the hash lookup causes the request to be sent to a server that is currently known to be unresponsive, a second hash lookup is performed that only takes the servers into account that are not known to be unresponsive. This is done to maximize the likelihood that all requests will end on the same server. If all configured servers are known to be unresponsive, the RADIUS proxy will fall back to the round-robin algorithm with the starting point determined by the first hash lookup to maximize the chance of getting any response to the request.

The following output displays a RADIUS server and policy configuration for servers referred from the RADIUS proxy.

```
config>service>vprn
   radius-server
      server "radius_server" address 10.100.100.2 secret "9OkclHYDDbo9eHrzFmuxiaO/
LAft3Pw"
                              hash2 port 1812 create
      exit
   exit

config>aaa
   radius-server-policy "radius_server_policy" create
      servers
         router 50
         access-algorithm hash-based
         source-address 10.1.1.1
         timeout min 1
         hold-down-time 2
         server 1 name "radius_server"
      exit
```

11.4.1.1.2 RADIUS Proxy — Cache Lookup

Local-user-database can be programmed to associate a host match with the RADIUS proxy cache instance. The host-match criterion is configurable, based on a subscriber attribute from the DHCP request.

The following output displays a RADIUS proxy cache lookup configuration.

```
config>subscriber-mgmt
   local-user-db "radius_ludb" create
       dhcp
           match-list service-id
           host "default" create
           auth-policy "auth_policy_1"
           match-radius-proxy-cache
              fail-action continue
              match mac
              server router 50 name "radius_proxy"
           exit
              no shutdown
```

```
        exit
        no shutdown
    exit
exit
```

If caching is enabled in the RADIUS proxy, then the actions on receiving DHCP message for the authenticated client includes the following:

- A host lookup is done in the local-user-database to find the RADIUS proxy cache for the subscriber.
- The field used to lookup the cache is configurable. It can include circuit-id or remote-id (present in sub-option in DHCP option-82), MAC@ or one of the other options in the DHCP packet. If a match is not found, the configured fail-action is executed. The default match field is MAC@. If the configured fail-action is “drop”, the DHCP DISCOVER is dropped. If the configured fail-action is “continue”, then the ESM host creation proceeds based on the authentication policy configured under the group-interface on which the DHCP packet is received.
- If a match is found, the parameters from original authentication accept in the cache are used to create the ESM host. If the group-interface is wlan-gw, then the ESM host is associated with the wlan-gw tunnel the (AP’s WAN IP@) and corresponding AP (MAC@ from the called-station-id in the authentication state).

11.4.1.1.3 RADIUS Proxy — Accounting

An ESM accounting-start is generated once the ESM host is created on successful authorization of DHCP against cached authentication state, and IP@ allocation is complete. The accounting-start contains information from locally cached 802.1x/EAP authentication such as calling-station-id, called-station-id, NAS-port-id, Subscriber-profile, SLA-profile, NAT port range for subscriber-aware NAT and so on.

If RADIUS proxy is configured as an accounting proxy in addition to authentication proxy, then the RADIUS proxy transparently forwards the accounting messages to the authentication server(s) referred from the RADIUS proxy, and can also load-balance. If caching is enabled, then the proxy can be configured to also track and locally act on the accounting messages, while still transparently forwarding these messages. The possible actions if accounting messages are tracked include the following:

- Accounting-start — The WIFI AP RADIUS client generates an accounting-start when a UE has successfully authenticated and associated with the AP. In cases where after mobility, the new AP does not re-authenticate due to key caching. accounting-start can be used as a mobility trigger on the WLAN-GW. Also, in cases where a UE associates with a single AP but pre-authenticates with

multiple APs in range, tracking mobility based on authentication can falsely associate a UE with incorrect AP. Mobility tracking based on authentication can be disabled via CLI (no track-authentication under radius-proxy cache), and instead be performed based on accounting-start. On receiving accounting-start, the RADIUS proxy on WLAN-GW finds the corresponding ESM host based on the calling-station-id attribute (typically the MAC@) of the subscriber) in accounting-start and associates the UE with the RADIUS client (for example, WIFI AP).

- Accounting-stop — The WIFI AP RADIUS client generates an accounting stop if it detects the UE has disassociated or is deleted due to inactivity or session timeout. The RADIUS proxy finds the corresponding ESM host based on the calling-station-id (typically the MAC@) of the subscriber. Note that if the called-station-id is filled out this must also match with what is currently stored as a security measure. When a UE moves the called-station-id should get updated and as such an accounting-stop from a previous AP cannot delete this UE anymore.
- The ESM host is deleted, an ESM accounting-sop message is sent, and the accounting-stop message from the AP is forwarded to the accounting-server.
- Accounting-ON or Accounting-OFF — This would be received from the AP if the AP has restarted. The RADIUS proxy will find all the impacted subscribers for the AP based on the called-station-id attribute (the AP's MAC@) in the accounting message, and delete all the corresponding ESM hosts.
- Interim Accounting Updates — If the client moves and re-associates with a new AP, the RADIUS client in the new AP generates interim-update. The RADIUS-proxy will locate the impacted ESM host, and update its state to point to the new AP's MAC@ (as available in called-station-id in the accounting message). The ESM interim-updates to accounting servers are sent on scheduled interval configured in accounting-policy, but with the updated information from the interim updates received from the AP.

11.4.2 Portal Authentication

For SSIDs without 802.11i/WPA2-based key exchange and encryption, it is common to authenticate the user by directing user's HTTP traffic to a portal, where the user is prompted for its credentials, which are verified against a subscriber database. The backend can optionally remember the MAC@ and subscriber credentials for a set period of time such that subsequent logins of the user do not require portal redirection. Some UEs support a client application (aka WISPr client), which automatically posts subscriber credentials on redirect, and parse HTTP success or failure response from the portal sever.

7750 WLAN-GW uses existing http-redirect action in IP filter to trigger redirect port-80 traffic. In case of open SSID, on receiving DHCP DISCOVER, MAC based authentication is performed with the RADIUS server as per configured authentication policy. The SLA-profile returned from RADIUS server in authentication-accept (or the default SLA-profile) contains the filter with http-redirect. Redirect via HTTP 302 message to the UE is triggered from the CPM. Once the user posts its credentials, RADIUS server generates a CoA-request message removing the http-redirect by specifying an SLA-profile without redirect action. If the portal authentication fails, the RADIUS server generates a disconnect-request message to remove the ESM host. In case of wlan-gw tunnel from the AP, the DHCP messages and data are both tunneled to the WLAN-GW. See Figure 179.

Figure 179 **Portal Authentication for Open SSIDs**

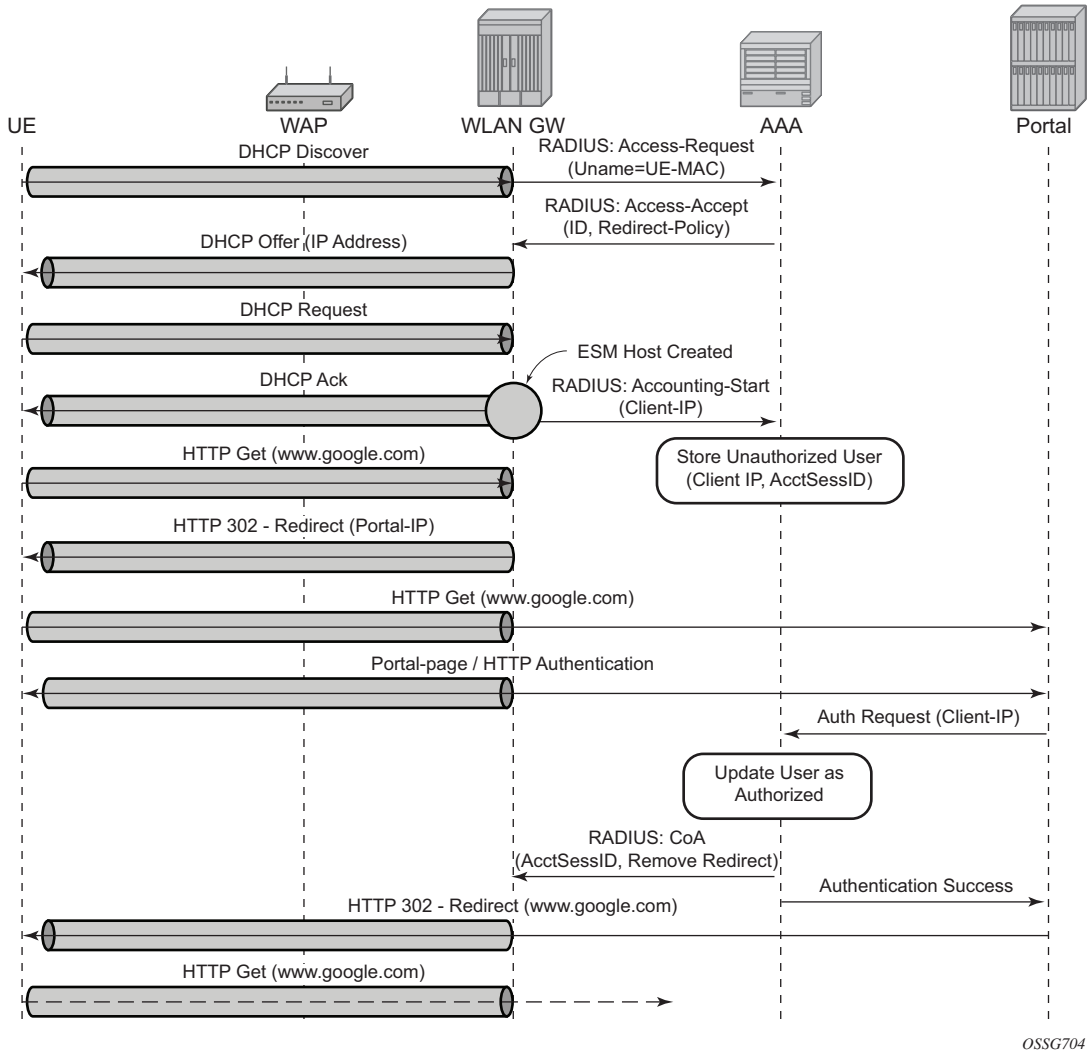

The following output displays a portal authentication for open SSIDs configuration example.

```
config>subscriber-mgmt
      sla-profile "portal-redirect" create
          ingress
             ip-filter 10
          exit
      exit
   exit

system>config>filter
   ip-filter 10 create
       entry 1 create
            match protocol udp
                dst-port range 67 68
            exit
            action forward
       exit
       entry 2 create
           match protocol tcp
               dst-port eq 80
           exit
           action http-redirect "http://www.google.ca"
       exit
   exit
exit
```

It is possible to view the subscriber HTTP redirect statistics by using the **show service** *id* **subscriber-hosts statistics** command. The statistics are captured per host and supports both IPv4 and IPv6. This command is only supported from CPM5 and up and SR-e platforms.

11.4.3 AA-based Portal Redirection

An application profile can be provided using RADIUS to be used during portal authentication. The WLAN-GW ISA redirects all traffic to a linked AA ISA, where all AA functionality can be used to enable a more extensive portal redirect. The WLAN-GW redirect and filter are completely bypassed, but the UE is kept in a portal authentication state. The UE can subsequently be promoted to DSM or ESM with a CoA as with regular portal functionality. When promoting to DSM, AA will be disabled unless a DSM application profile is provisioned using configuration or RADIUS.

See Application Assurance Support for more information on enabling AA functionality for a WLAN-GW.

11.5 Address Assignment

The address to the UEs can be assigned via local DHCP server from locally defined pools, or from RADIUS server via local DHCP proxy, or from an external DHCP server. Subscriber-interface and group-interface are configured as part of normal ESM configuration. In case of wlan-gw, the group interface is wlan-gw enabled. Subnets on the subscriber interface are used for the pools from which the DHCP local server assigns addresses to UEs.

The following output displays an address assignment configuration example.

```
config>service>vprn
    dhcp
       local-dhcp-server "dhcp" create    #### create local DHCP server
           pool "1" create                #### define Pool
               options
                   dns-server 10.8.8.8 8.8.4.4
                   lease-time min 5
               exit
               subnet 203.0.113.255/30 create
                  options
                     subnet-mask 239.255.0.0
                     default-router 10.203.254.181
                  exit
                  address-range 128.203.254.182 128.203.254.183
               exit
             exit
         exit
     exit

    interface "DHCP-lb" create          #### loopback interface with DHCP server
        address 10.1.1.1/32
        local-dhcp-server "dhcp"
        loopback
    exit

    subscriber-interface "sub-int" create          #### subscriber interface
         address 172.31.235.235/30                #### Subnets out of which UE
         address 172.31.245.245/16                ######  addresses are allocated.
         group-interface "group-int" wlgw create
             sap-parameters
                 sub-sla-mgmt
                   def-sla-profile "sla_def"
                   def-sub-profile "sub_def"
                   sub-ident-policy "sub_ident"
                 exit
              exit
         exit
         dhcp
             proxy-server
                 emulated-server 10.10.0.1   #### proxy to get IP address from AAA
                 lease-time min 5           #### or from DHCP server. Can provide
                 no shutdown        #### split lease (shorter lease towards client,
              exit                 #### and longer lease towards AAA or DHCP server.
              no option
```

```
    server 10.1.1.1                ##### DHCP local server
        trusted
        lease-populate 32000
        gi-address 172.31.245.245
        user-db "radius_ludb"      #### LUDB for proxy cache co-relation
        no shutdown
     exit
exit
```

11.6 WIFI Mobility Anchor

7750 WLAN-GW supports seamless handling for UE mobility, when a UE moves from one AP to another, where the new AP is broadcasting the same SSID, and is anchored on the same WLAN-GW. In case of open SSID, when the UE re-associates with the same SSID on the new AP and already has an IP@ from association with previous AP, the UE can continue to send and receive data. The WLAN-GW learns the association of the UE's MAC address to the GRE tunnel corresponding to the new AP, and updates its state on the MS-ISA as well as on the CPM. The UE continues to be anchored on the same anchor MS-ISA, thereby avoiding any disruption in ESM functions (SLA enforcement and accounting). State update based on data learning results in fast convergence after mobility and minimal packet loss. The data-triggered mobility can be turned on via configuration. Mobility trigger can be configured to be restricted to special Ethernet IAPP frame (originated by the AP with the source MAC of UE).

For 802.1x/EAP based SSIDs, by default the AP requires re-authentication to learn the new session keys (PMK). 7750 SR as WLAN-GW RADIUS proxy infers mobility from the re-authentication, and updates the ESM host to point to the new AP. The new AP's IP address is derived from the RADIUS attribute NAS-IP-address. The re-authentication also provides the new session keys to the AP in access-accept RADIUS response. In case the WIFI AP or ACs are capable of PMK key caching or standard 802.11r (or OKC, the opportunistic key caching pre-802.11r), the re-authentication on re-association can be avoided. In this case the UE can continue to send data, and the WLAN-GW can provide fast data-triggered mobility as defined in context of open SSIDs.

The following output provides a mobility anchor configuration example.

```
config>service>ies>
config>service>vprn>
subscriber-interface <if-name>
group-interface <if-name> wlangw
wlan-gw
     [no] router (base | <vprn-id>) # tunnel service context
     [no] wlan-gw-group <group-id>
      ....snip
       mobility
         [no] trigger {data | iapp}
         [no] hold-time <seconds> // [0..255 secs]
       exit
     exit
exit
```

11.7 Wholesale

With EAP the AAA server can look at the realm from the user credential (IMSI) in authentication request and appropriately provide the service context in retail-service-id, for the ESM host corresponding to the UE.

For open SSID, the decision can be made by the AAA server based on the SSID, as learned by any of the methods described in Authentication.

The SSID is passed to the AAA server in initial MAC based authentication on DHCP DISCOVER. The retail-service-id can be returned in access-accept. This assumes the AP broadcasts unique SSID per retail provider, and inserts it in Option82 as a DHCP relay-agent. As an alternative to SSID in option-82, the AP can insert a unique dot1Q tag per retail provider, before tunneling the Ethernet frame, using single GRE tunnel per AP to the WLAN-GW. 7750 supports configuring a map of .dot1Q tags to retail-service-id. Therefore, the determination of the retail provider for the subscriber can be made in the data plane when DHCP is received, and the subscriber state can be created and processed in the right service context.

The following output displays a wholesale configuration example.

```
config>service>ies>
config>service>vprn>
subscriber-interface <if-name>
group-interface <if-name> wlangw
wlan-gw
[no] router (base | <vprn-id>) # tunnel service context
[no] wlan-gw-group <group-id>
....snip
vlan-tag-ranges # Precedence for retail-service-id:
# RADIUS, vlan-retail-service-map, default-retail-svc
  [no] vlan start <start-tag> end <end-tag> retail-svc-id <svc-id>
  [no] default-retail-svc-id
exit
exit
exit
```

11.8 CGN on WLAN-GW

Both LSN and L2-aware NAT for WIFI subscribers over wlan-gw tunnels is supported. NAT on WLAN-GW is only supported for locally terminated subscribers and not for GTP tunneled subscribers. NAT can be performed on the same set of ISAs that are used for WLAN-GW functions, by referring to the WLAN-GW ISA group from NAT configuration. Alternatively, dedicated set of ISAs can be used for NAT function by creating and referencing a separate NAT-group. Configuration related to LSN and L2-aware NAT is provided in the SR OS Multiservice Integrated Service Adapter Guide.

11.9 Lawful Intercept on WLAN-GW

Mirroring traffic for WIFI subscribers to a mediation device, when the subscriber is under legal intercept is supported. The mirroring function is performed on the anchor IOM where the subscriber is anchored. Both Ether and IP-only mirror is supported. With Ether mirror, VLAN tags which are part of internal SAP between ISA and IOM, are included in the mirrored Ethernet frame of the subscriber. IP-only mirror includes the IP header and the payload. Conventional IP-only mirror service can be used with direct p2p or MPLS (for remote mirroring) connection to the mediation device. Routable-encapsulation is also supported. Both IP/UDP encapsulation with optional shim-header for subscriber correlation on the mediation device, and IP/GRE encapsulation is supported with routable-encapsulation of mirrored data. LI can be triggered via CLI, SNMPv3 or RADIUS, as supported with ESM. RADIUS triggered LI can be via LI related VSAs in access-accept or in CoA. The CoA is keyed on accounting-session-id. LI is supported for both local and GTP tunneled subscribers.

Existing LI support with ESM is described in the *7450 ESS, 7750 SR, 7950 XRS, and VSR OAM and Diagnostics Guide*.

11.10 WLAN Location Enhancements

This feature adds configurable support for learning and reporting the AP's MAC address (which represents WLAN location of the UE), to the AAA server. Support is also added for triggered interim accounting-updates to report the AP's MAC@ to the AAA server.

11.10.1 Triggered Interim Accounting-Updates

Using location based policy for WIFI subscribers is important. The business logic in AAA could use the location of the subscriber. Therefore, it is important to notify location change of the subscriber to AAA. Standard way to do this is by generating an interim accounting update when the WLAN-GW learns of the location change for a subscriber. The location for a WIFI subscriber can be inferred from MAC@ (preferred) or WAN IP@ of the AP.

For open-SSID, learning about mobility could be data-triggered or IAPP packet triggered. If triggered, interim accounting-update is configured via CLI, then on detecting a location change for the UE, an interim accounting-update is sent immediately to the AAA server with the new AP's MAC@ (if already known to WLAN-GW). The accounting-update contains NASP-port-id (which contains the AP's IP@), and circuit-id (from DHCP option-82) which contains AP's MAC@ and SSID. In case of data-triggered mobility, if the new AP's MAC@ is not already known to WLAN-GW, a GRE encapsulated ARP packet is generated towards the AP to learn the MAC@ of the AP. The AP is expected to reply with a GRE encapsulated ARP response containing its MAC@. The generation of ARP to learn the AP's MAC@ is controlled via CLI. The GRE encapsulated ARP packet is shown in Figure 180.

Figure 180 **GRE Encapsulated ARP Request**

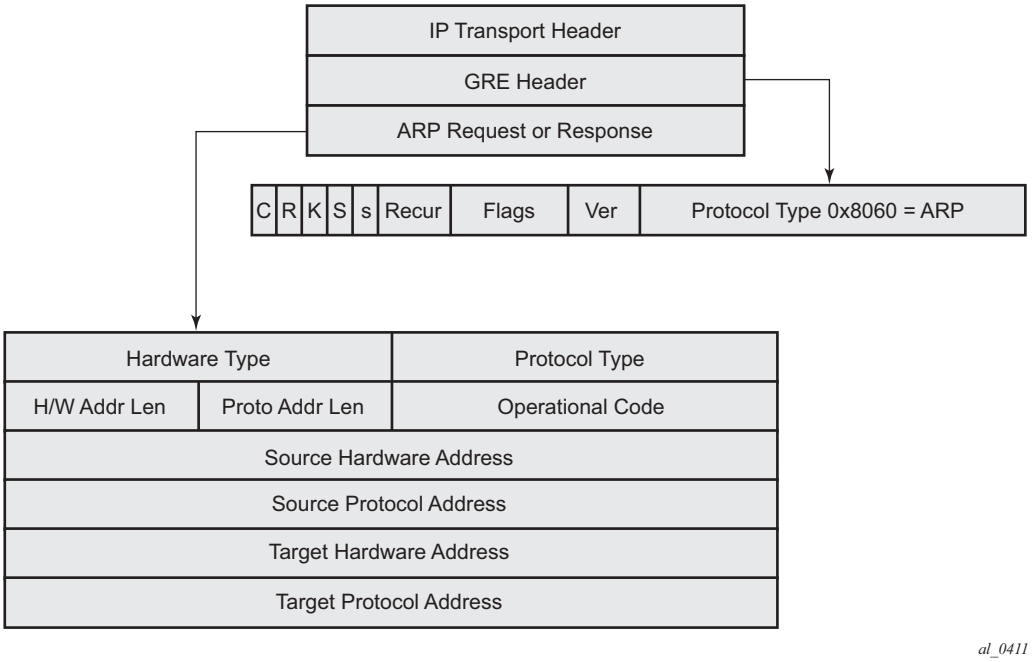

The standard ARP request must be formatted as follows:

- Hardware Type = Ethernet (1)
- Protocol Type = 0x0800 (IPv4)
- H/W Addr Len = 6
- Proto Addr Len = 4
- Operational code (1 = request)
- Source hardware address = WLAN-GW MAC@
- Source protocol address = Tunnel endpoint IP@ on WLAN-GW
- Target hardware address = Unknown
- Target protocol address = WAN IP@ of the AP (source IP in GRE packet)

The AP must generate a GRE encapsulated ARP response when it receives the GRE encapsulated ARP request for its WAN IP@ (that is used to source tunneled packets). The standard ARP response should be formatted as follows:

- Hardware Type = Ethernet (1)
- Protocol Type = 0x0800 (IPv4)
- H/W Addr Len = 6
- Proto Addr Len = 4

- Operational code (2 = response)
- Source hardware address = AP MAC@
- Source protocol address = WAN IP@ of AP (used for sourcing tunneled packets)
- Target hardware address = source hardware address from the request
- Target protocol address = source protocol address from ARP request

For 802.1x/EAP SSID, the location change (mobility) is learned from an interim-accounting update from the AP. The called-station-Id (containing the AP MAC@) is compared against the current stored called-station-Id that the subscriber is associated with. If the called-station-id is different than the received interim accounting update is immediately forwarded to the accounting server, if triggered interim accounting-update is configured via CLI. In previous releases, the interim-update received from the AP is not immediately forwarded by the accounting proxy. Only a regularly scheduled interim-update is sent.

11.10.2 Mobility Triggered Interim Updates with Counters

This feature supports configurable inclusion of counters for an ESM host (UE) in triggered interim updates, due to mobility, generated from WLAN-GW. Triggered interim updates provide a way for RADIUS to learn of the current location of the UE (such as AP MAC).

The feature also supports a configurable hold down time to protect against very frequent mobility events. The CPM must communicate with the IOM for each mobility event to obtain the counters. By default, when inclusion of counters is enabled, no hold-down time is imposed on mobility triggered interim updates. If a hold down time is configured, the first mobility event triggers an interim update with counters included and start the hold down timer. As long as the hold down timer has not expired, interim updates due to a mobility event will not be sent.

The inclusion of counters and hold down time are applicable to all access-types (soft-GRE, soft-L2TPv3, L2-AP, SAP, PW-SAP, and so on), and for DHCP, data-triggered, and EAP authentication or re-authentication triggered mobility.

The configuration for counter inclusion subject to optional hold down time is only applicable to ESM. For counters to be included in the mobility-triggered interim updates, the general **std-acct-attributes** or **detailed-acct-attributes** in the include-attribute in the **radius-accounting-policy** must also be enabled. With DSM, counters are always included in mobility-triggered interim updates as long as the **isa-radius-policy** has inclusion for either or both **frame-counters** or **octet-counters** enabled. Unlike ESM, no new WLAN-GW-specific configuration is required for the counter inclusion.

The feature is supported at the maximum ESM host scale.

```
config>router>wlan-gw>mobility-triggered-acct>
     interim-update
     interim-update include-counters [hold-down <seconds>]
     no interim-update
<include-counters>   : keyword
<seconds>            : [60..86400]
```

11.10.3 Operational Support

Following command shows if GRE encapsulated ARP request is enabled.

```
*A:Dut-C# show router 4 interface "grp-vprn_ue-2/1/2:50" detail
===============================================================================
Interface Table (Service: 4)
===============================================================================
-------------------------------------------------------------------------------
Interface
-------------------------------------------------------------------------------
If Name          : grp-vprn_ue-2/1/2:50
Sub If Name      : ies-4-20.0.0.1
Red If Name      :
Admin State      : Up                  Oper (v4/v6)      : Up/Up
Protocols        : None

WLAN Gateway details
Administrative state        : in-service
Router                      : 50
IP address                  : 10.1.1.3
IPv6 address                : 2001:db8::0
ISA group ID                : 1
Egr shaping                 : none
Egr shape multi UE only     : false
Egr qos policy ID           : (Not Specified)
Egr scheduler policy        : (Not Specified)
Egr agg rate limit (kbps)   : (Not Specified)
Egr qos resrc hold time (s) : 0
Mobility trigger            : data iapp
Mobility ARP AP             : enabled
Mobility hold time (s)      : 0
Default retailer service    : (Not Specified)
TCP MSS adjust              : (Not Specified)
Number of tunnels           : 0
Last management change      : 02/19/2014 17:48:52
```

11.11 3G/4G Interworking

The WLAN-GW supports 3G/4G interworking based on a per-UE GTP tunnel from WLAN-GW to the mobile packet core (GGSN or P-GW). For more details on the GTP uplink setup, see GTP. The bridged WiFi AP connectivity with the WLAN-GW can be WLAN-GW-based (soft-tunnel/l2-ap), or it can be a native Layer 2 (regular group-interface).

11.11.1 Signaling Call Flow

The decision to setup a GTP tunnel for a subscriber or locally breakout subscriber's traffic is AAA based, and received in authentication response. If the traffic is to be tunneled to the PGW or GGSN, the signaling interface or PGW/GGSN interface would be provided via AAA. Absence of these attributes in the authentication response implicitly signifies local-breakout.

11.11.1.1 GTP Setup with EAP Authentication

Once the EAP authentication completes as described in the section on authentication, the RADIUS proxy caches the authentication response, including any attributes related to GTP signaling. Subsequently DHCP is initiated from the UE. On receiving DHCP DISCOVER, the RADIUS proxy cache is matched to get the AAA parameters related to the UE from the original authentication response. If PGW/GGSN (mobile gateway) IP address is not present in cached authentication, DNS resolution as described in section 1.2 is initiated for the WLAN APN obtained from AAA (in the cache) or for locally configured APN in the service associated with the UE. The DNS resolution provides a set of IP addresses for the mobile gateways. The GTP tunnel setup is attempted to the selected mobile gateway. The IP address provided by PGW/GGSN in the GTP response is returned in DHCP offer to the UE. The WLAN-GW acts as a DHCP to GTP proxy. The WLAN-GW is the default-GW for the UE. Any packets from the UE are then GTP tunneled to the mobile gateway. If the UE requests an IP address (for which it may have an existing lease on one of its interface) via DHCP option 50 in the DHCP request, then WLAN-GW sets the "handover bit" in the GTP session create message, and indicates the requested address in the PDN Address Allocation (PAA) field. This allows the PGW to look for existing session corresponding to the signaled IMSI and APN (with potentially different RAT-Type) and return its existing IP address in session create response. The old session and bearer is deleted by the PGW. The signaling of "handover bit" is supported with S2a and S2b (Release 10.0 and beyond). The IP address cannot be preserved over the Gn interface. The call flow in Figure 184 shows basic GTP

setup (with S2a), the output provided on Figure 184 shows IP address preservation across inter-access (WIFI <-> 4G) moves.

DHCP release or lease timeout on WLAN-GW will result in deletion of the GTP tunnel corresponding to the UE. The session or PDP context deactivation from PGW/GGSN will also result in removal of the GTP state for the UE and the corresponding ESM host on WLAN-GW. In this SR OS release, only default bearer (or primary PDP context) for single default APN is handled over WIFI. GTP path-management messages (echo request and reply) are supported. Mandatory IEs are supported in GTP signaling. Hard coded default values are signaled for QoS and charging related IEs. For GTPv2, the bearer is signaled as non-GBR bearer with QCI value of 8, and MBR/GBR values of 0. APN-AMBR default values signaled are 20 Mb/s / 10 Mb/s downstream/upstream. For GTPv1, reliability and priority classes default to “best-effort”, allocation/retention priority defaults to 1, and the default peak-rate corresponds to class 9 (bit-wise 1001) which is slightly over 2 Mb/s. Charging characteristics IE which contains a 16 bit flag defaults to 0. In the future, RADIUS returned values or locally configurable values is signaled in QoS and charging IEs.

The IP address is returned in the create PDP context response or Create session response. The DNS server addresses for the UE are returned in IP control protocol (IPCP) option in a PCO IE in the response. The default gateway address provided to the UE in DHCP is auto-generated algorithmically on the WLAN-GW from the IP address returned by the PGW/GGSN for the UE. The WIFI AP is required to provide a split-horizon function, where there is no local switching on the AP, and all communication to/from any AP is via WLAN-GW. The WLAN-GW implements proxy-ARP and forwards all received traffic from the UE into the GTP tunnel. In the future, the default-GW address to be returned to the UE could be obtained in a PCO from the PGW/GGSN. The GTP-U processing of data packets is done in the IOM.

11.11.2 Location Notification in S2a

This feature adds support on WLAN-GW for reporting UE’s WLAN location (TWAN Identifier IE) and cellular location (ULI IE) over S2a interface to PGW and UE’s cellular location (ULI IE) to GGSN (over Gn interface). Location information is useful for charging on PGW/GGSN.

11.11.2.1 WLAN Location over S2a

The WLAN location information consists of the *TWAN Identifier IE* as described in 29 274 V11.6.0 (2013-04) and is sent in GTPv2 “create session request” message. If present, this IE carries BSSID (MAC address of the AP) and the SSID. WLAN-GW learns the AP’s MAC@ from calling-station-id attribute in the RADIUS messages from the AP (both authentication and accounting messages) or from circuit-id in DHCP DISCOVER or REQUEST messages. In this release, the IE is only sent at session creation time. Therefore, it reports location on initial attach, on handover from LTE to WIFI, and on AP mobility across WLAN-GWs. Mobility across APs anchored on the same WLAN-GW does not result in location update. 3GPP release 11 does not define location update mechanism for S2a.

By default, location is not reported. It can be enabled with CLI.

```
config>subscr-mgmt>gtp
   peer-profile "pgw-west-mn01"
      [no] report-wlan-location
```

11.11.2.2 Cellular Location over S2a

The “User Location Info” IE is included in “Create Session Request and is described in 3GPP TS 29.274 version 8.1.1 Release 8. The encoding for individual location identifiers (CGI, SAI, RAI, TAI, and ECGI) is also defined in the same reference (as shown in Figure 181).

Figure 181 **User Location Information**

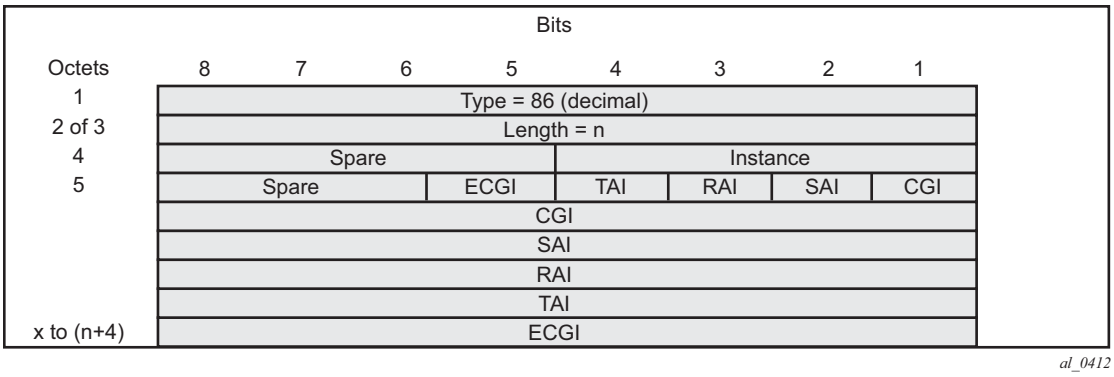

The AP's MAC@ and IP@ are provided to AAA server in RADIUS messages during EAP authentication and accounting. If AAA provides the cellular location (corresponding to this AP) in 3GPP attribute *3GPP-User-Location-Info* in access-accept, and location reporting is enabled via CLI. The ULI IE is included in GTPv2 "create session request". The *3GPP-User-Location-Info* attribute is described in 3GPP TS 29.061 v9.3.0.

11.11.2.3 Cellular Location over Gn Interface

The "User Location Info" IE (as shown in Figure 182) can be included in create-pdp-context message as described in 3GPP TS 29.060 V10.1.0. The geographic location type field describes the type of location included in the "Geographic Location" field that follows. The location can be CGI (cell global identification), SAI (service area identity), or RAI (routing area identity). The formats for these location identifiers are defined in the same reference *3GPP TS 29.060 V10.1.0*.

Figure 182 **User Location Information IE**

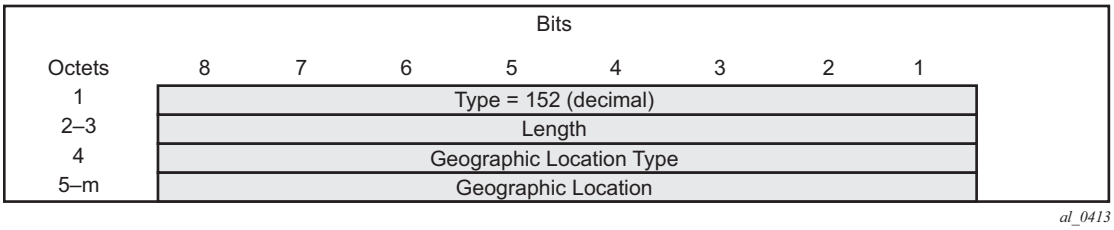

AP MAC address and SSID is reported to AAA (including changes on mobility). AAA can then specify the ULI IE contents based on static mapping of AP's MAC address to one of the cellular location types (CGI, SAI or RAI). AAA should provide the cellular location in 3GPP attribute 3GPP-User-Location-Info (below) in access-accept. The attribute is described in 3GPP TS 29.061 v9.3.0.

In case a UE moves to a different WLAN-GW, UE is authenticated based on data-trigger. In this case, the AAA server can provide the WLAN location (AP's MAC@ and SSID) in called-station-ID attribute and cellular location in 3GPP-User-Location-Info attribute. The WLAN location is then encoded in TWAN identifier in "create session request" message, and the cellular location is encoded in the ULI IE.

11.11.2.4 Operational Support

The following command shows state of location reporting (enabled/disabled).

```
*A:Dut-C>config>subscr-mgmt>gtp>peer-profile$ /show subscriber-mgmt gtp
```

```
peer-profile "test"

===============================================================================
GTP peer profile "test"
===============================================================================
Description                 : (Not Specified)
Retransmit timeout (s)      : 5
Retransmit retries          : 3
Keepalive interval (s)      : 60
Keepalive retries           : 4
Keepalive retry timeout (s) : 5
Time to live                : 255
Interface type              : s2a
Charging char home UE       : (None)
Charging char roaming UE    : (None)
Session hold time (s)       : 30
Report WLAN location        : enabled
Last management change      : 02/21/2014 16:31:12
GGSN uplink GBR (Kbps)      : 5000
GGSN uplink MBR (Kbps)      : 5000
GGSN downlink GBR (Kbps)    : 2000
GGSN downlink MBR (Kbps)    : 2000
GGSN Alloc/Retention Prio   : 1
GGSN last management change : 02/19/2014 17:31:55
PGW uplink GBR (Kbps)       : 0
PGW uplink MBR (Kbps)       : 0
PGW downlink GBR (Kbps)     : 0
PGW downlink MBR (Kbps)     : 0
PGW Alloc/Retention Prio    : 1
PGW Qos Class ID            : 8
PGW last management change  : 02/19/2014 17:31:55
===============================================================================
```

11.12 Migrant User Support

"Migrant users" are UEs that connect to an SSID, but move out of the range of the access-point before initiating or completing authentication. For open-SSIDs, a migrant user may stay in the range of the access-point just enough to get a DHCP lease from the WLAN-GW. In real WIFI deployments with portal authentication, it has been observed that a large percentage of users are migrant, such as get a DHCP lease but do not initiate or complete authentication. Prior to this feature, an ESM host is created when DHCP completes. This results in consumption of resources on both CPM and IOM, limiting the ESM scale and performance for fully authenticated active users. This feature adds support to only create an ESM host after a user has been fully authenticated, either via web portal or with a AAA server based on completing EAP exchange. In addition, with this feature L2-aware NAPT is enabled on the ISA, such that each UE gets the same shared configured inside IP@ from the ISA via DHCP. Until a user is authenticated, forwarding of user traffic is constrained (via policy) to only access DNS and portal servers. Each user is allocated a small number of configured NAT outside ports to minimize public IP address consumption for unauthenticated users. Once the user is successfully authenticated, as indicated via a RADIUS COA on successful portal authentication, an ESM host is created, and the L2-aware NAT is applied via a normal per-subscriber NAT policy. The inside IP address of the user does not change. The outside IP pool used is as per the NAT policy, and the L2-aware NAT could be 1:1 or NAPT with larger number of outside ports than in the un-authenticated phase. If a user is already pre-authenticated (for example, if RADIUS server remembers the MAC@ of the UE from previous successful portal authentication), then the initial access-accept from RADIUS will trigger the creation of the ESM host.

Migrant user support is only applicable to EAP based closed SSIDs when RADIUS-proxy is not enabled on WLAN-GW. This is described in Migrant User Support with EAP Authentication.

11.12.1 Migrant User Support with Portal-Authentication

11.12.1.1 DHCP

Based on DHCP and L2 NAT configuration on the ISA, IP address is assigned to the user via DHCP. A different DHCP lease-time can be configured for an unauthenticated user and an authenticated user for which an ESM host has been created. DHCP return options, for example, DNS and NBNS server addresses can be configured. This configuration can be per wlan-gw group interface or per VLAN

range (where a VLAN tag corresponds to an SSID). After the DHCP ACK is sent back to the UE from the ISA, the UE is created on the ISA in “migrant (or unauthenticated) state”. ARP requests coming from the UE in migrant state is responded to from the ISA. The authentication to RADIUS is triggered on receiving first L3 data-packet as opposed to on DHCP DISCOVER.

11.12.1.2 Authentication and Forwarding

The authentication is initiated from RADIUS client on the ISA anchoring the user, based on an isa-radius-policy (configured under aaa) and specified on the wlan-gw group-interface. The initial access-accept from RADIUS can indicate if a user needs to be portal authenticated or is a pre-authenticated user. The indication is based on inclusion of a “redirect policy” applicable to the user, in a VSA (Alc-Wlan-Portal-Redirect, type = string). The access-accept can also include a redirect URL VSA (Alc-Wlan-Portal-Url, type = string) for the user. An empty Alc-Wlan-Portal_redirect VSA forces the use of locally configured redirect policy. Also, if neither of the above two VSAs are included, then this indicates a “pre-authenticated user”, and an ESM host is created for the subscriber with subscriber-profile and other subscriber configuration from access-accept, and from here normal ESM based forwarding occurs for the subscriber.

However, if a user needs portal authentication (as indicated in access-accept), then while the user is pending authentication, forwarding is restricted to DNS and portal servers via the redirect policy. The redirect policy is an IP ACL that restricts forwarding based on IP destination, destination port, and protocol, and also specifies http-redirect for http traffic that does not match any of the forwarding rules. The URL for re-direct is configured in the redirect policy or can be provided in authentication-accept. A maximum of 16 redirect policies can be created in the system, with a maximum of 64 forward rules across all redirect policies. During this “authentication pending” phase all forwarded traffic is subjected to L2-aware NAT on the ISA. The NAT policy to use for these users can be configured on the wlan-gw interface or per VLAN range under the wlan-gw interface. After an access-accept has been received from RADIUS for such a user, the next http packet triggers a redirect function from the ISA, and an http 302 is sent to the client. The redirect can be configured to append original-URL, subscriber’s MAC address and IP address to the redirect URL sent back in http 302. The client presents its credentials to the portal and once it is successfully authenticated, a COA is generated from the RADIUS server (triggered by the portal). The COA message triggers creation of an ESM host with the subscriber configuration contained in the COA such as subscriber-profile, SLA-profile, NAT-profile and application-profile. From this point normal ESM based forwarding occurs for the subscriber.

See, Migrant User NAT Configuration for configuration information related to migrant users.

11.12.2 Migrant User Support with EAP Authentication

Migrant user support can only be used for closed SSIDs when there is no RADIUS-proxy configured on WLAN-GW. If no RADIUS proxy is configured, then initial RADIUS request carrying EAP from the AP is normally forwarded to a RADIUS server. The RADIUS exchange is between AP and the AAA server, and no information from EAP authentication is cached on the WLAN-GW. The subsequent DHCP DISCOVER after successful EAP authentication is received on the ISA. However, for subscriber that needs to be GTP tunneled to PGW/GGSN, the DHCP is forwarded to the CPM, where it triggers a RADIUS authorization. RADIUS correlates the MAC address with calling-station-id from EAP authentication for the user. GTP tunnel initiation, and ESM host creation then follows after receiving an access-accept. However, for a “local-breakout” subscriber DHCP and L2-aware NAT is handled on the ISA (as in the case for migrant users with portal based authentication). Shared inside IP address can be handed out to each subscriber. The first L3 packet triggers MAC address based RADIUS authorization from the ISA. RADIUS server can correlate the EAP authentication with the MAC address of the user and send an access-accept. This triggers ESM host creation as normal.

For closed SSIDs with EAP authentication, if a RADIUS proxy function is configured on WLAN-GW, then the initial EAP authentication from the AP is processed by the RADIUS-proxy on the CPM, and is forwarded to the RADIUS server based on configured authentication policy. Based on authentication response, ESM host creation with local DHCP address assignment or GTP tunnel initiation proceeds as usual.

11.12.3 Data Triggered Subscriber Creation

With **data-triggered-ue-creation** configured under wlan-gw group interface or per VLAN range (such as, per one or more SSIDs), the first UDP or TCP packet received on WLAN-GW ISA from an unknown subscriber (with no prior state, such as an unknown MAC address) will trigger RADIUS authentication from the ISA. The authentication is based on configured isa-radius-policy (under aaa context). If RADIUS authentication succeeds, then ESM host is created from the CPM. The ESM host can get deleted based on idle-timeout. Data-triggered authentication and subscriber creation enables stateless inter WLAN-GW redundancy, as shown in Figure 183. If the AP is configured with a backup WLAN-GW address (or FQDN), it can tunnel subscriber traffic to the backup WLAN-GW, when it detects failure of the primary WLAN-GW (based on periodic liveness detection). With “data-triggered-ue-creation” configured, the first data packet results in authentication and ESM host creation on the backup WLAN-GW. If the subscriber had obtained an IP address via DHCP with L2-aware NAT on the primary WLAN-GW, it can retain it with L2 aware NAT on the backup WLAN-GW. The NAT outside pool for the subscriber changes on

the backup WLAN-GW based on local configuration. For a subscriber that needs to be anchored on GGSN/PGW (as indicated via RADIUS access-accept), RADIUS server returns the IP address of PGW/GGSN where the UE was anchored before the switch-over. GTP tunnel is then signaled with “handover indication” set. The PGW/GGSN must return the requested IP address of the UE, which is the address with which the UE originated data packet that triggered authentication.

The same data-triggered authentication and subscriber creation is also used to support inter WLAN-GW mobility, such as when a UE moves form one AP to another AP such that the new AP is anchored on a different WLAN-GW. This is shown in Figure 184.

Figure 183 **N:1 WLAN-GW Redundancy Based on Data-Triggered Authentication and Subscriber Creation**

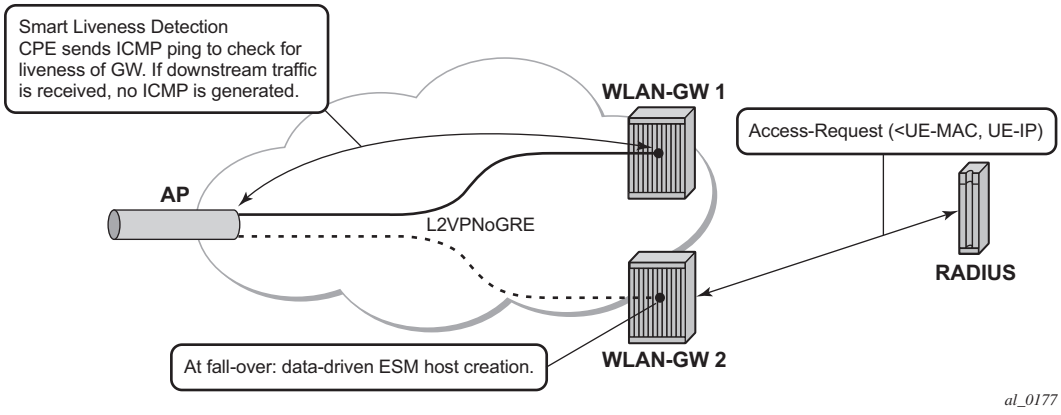

Figure 184 **Inter WLAN-GW Mobility Based on Data-Triggered Authentication and Subscriber Creation**

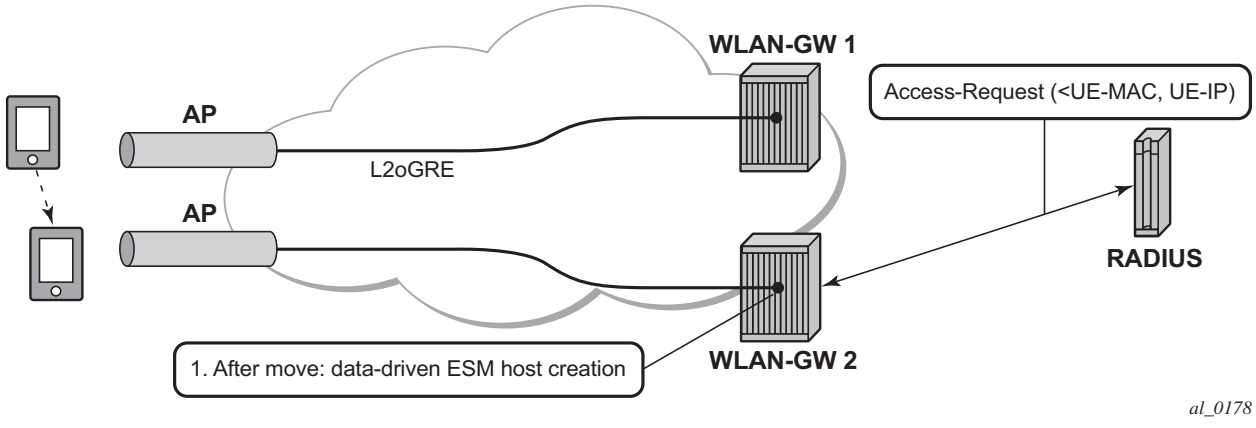

The following output displays the configuration for migrant user support and “data-triggered” subscriber creation.

Migrant User NAT Configuration

```
#--------------------------------------------------
 NAT configuration for migrant and authenticated users
#--------------------------------------------------
service

  vprn 300 customer 1 create

     nat
       inside
          l2-aware
               address 10.20.12/16
          exit
       exit
       outside
           pool "migrant_outside_pool" nat-group 1 type wlan-gw-anchor create
                address-range 10.22.0.0 10.22.0.255 create
                exit
                no shutdown
           exit
           pool "wifi_outside_pool" nat-group 1 type l2-aware create
                address-range 10.0.0.0 10.0.0.255 create
                exit
                no shutdown
           exit
       exit
     exit
  exit

  nat
   nat-policy "migrant_nat_300" create
        pool "migrant_outside_pool" router 300
        timeouts
             tcp-established min 1
        exit
   exit

   nat-policy "wifi_nat_300" create
        pool "wifi_outside_pool" router 300
   exit

 exit


#------------------------------------------------------------------------------
echo "AAA Configuration" - ISA-RADIUS-Policy for authentication from WLAN-GW ISA
#------------------------------------------------------------------------------
    aaa
        isa-radius-policy "wifi_isa_radius" create
            description "Default authentication policy for migrant users"
            password "i2KzVe9XPxgy4KN2UEIf6jKeMT3X4mT6JcUmnnPZIrw" hash2
            servers
                router "Base"
                source-address-range 10.100.100.4
                server 1 create
                    authentication
                    coa
                    ip-address 10.100.100.2
```

```
                    secret "ABIQRobhHXzq13ycwqS74FSrj.OdTwh5IdjhRB.yAF." hash2
                    no shutdown
                exit
            exit
        exit
        radius-server-policy "radius_server_policy" create
            servers
                router "Base"
                server 1 name "radius_server"
            exit
        exit
    exit

#--------------------------------------------------
echo "Subscriber-mgmt Configuration" - Redirect Policy
#--------------------------------------------------
    subscriber-mgmt
        http-redirect-policy "migrant_redirect" create
            url "portal.ipdtest.nokia.com:8081/start/?mac=$MAC&url=$URL&ip=$IP"
            portal-hold-time 10
            forward-entries
                dst-ip 10.8.8.1 protocol tcp dst-port 8081
                dst-ip 10.8.8.7 protocol tcp dst-port 8007
                dst-ip 10.8.8.8 protocol udp dst-port 53
            exit
        exit
     exit
service

#------------------------------------------------------------------
echo "migrant user configuration under wlan-gw group interface"
#------------------------------------------------------------------

  vprn 300 customer 1 create

    subscriber-interface "ies-4-20.10.1.1" create
        address 10.20.12/16

        group-interface "grp-vprn_ue-2/1/2:51" wlangw create
            sap-parameters
                sub-sla-mgmt
                    def-sla-profile "slaprof_1"
                    def-sub-profile "subprof_1"
                    sub-ident-policy "identprof"
                exit
            exit
            dhcp
                proxy-server
                    emulated-server 10.20.12.12
                    no shutdown
                exit
                trusted
                lease-populate 32767
                user-db "radius_ludb"
                no shutdown
            exit
            host-connectivity-verify interval 1000
            wlan-gw
                gw-address 10.1.1.4
```

```
                    mobility
                        hold-time 0
                        trigger data iapp
                    exit
                    router 50
                    wlan-gw-group 1
                    vlan-tag-ranges
                         range start 100 end 100
                              authentication
                                   authentication-policy "wifi_isa_radius"
                              exit
                              data-triggered-ue-creation
                              dhcp
                                 l2-aware-ip-address 10.1.1.2
                                 primary-dns 10.1.1.1
                                 secondary-dns 10.1.1.1
                                 no shutdown
                              exit
                              nat-policy "migrant_nat_4"
                          exit
    exit
                     no shutdown
                exit
          exit
     exit
exit
```

11.13 Distributed Subscriber Management (DSM)

With this feature, once the UE is successfully authenticated (portal, auto-signed-in, or EAP), the corresponding subscriber can be created on the anchor ISA, and both control plane and forwarding plane for the subscriber are handled on the ISA. This mode of subscriber management is henceforth referred to as **Distributed Subscriber Management** (DSM).

Prior to this feature, only ESM is supported for WLAN UEs, where the ESM host state is created on the IOM/IMMs from the CPM (triggered by the ISA on successful authentication). With ESM, the initial DHCP process and authentication could be triggered from the ISA (based on a per VLAN-range configuration for DHCP) under the group-interface with of type **wlangw**. However, control plane operations after the ESM host creation (such as accounting and DHCP renews) are handled on the CPM.

With DSM, in addition to initial DHCP and authentication, once the subscriber state exists on the anchor ISA, accounting and DHCP renews are also handled from the anchor ISA for the UE. This allows a higher UE scale and better control plane performance (including DHCP transactions per second, rate of authentications, and web redirects) due to load-balancing amongst set of ISAs in the WLAN-GW group. With DSM, the UE data-plane functions (such as per UE IP filtering, ingress/egress policing, legal intercept, per UE counters, and web-redirect) are performed on the ISA.

The decision to create an authenticated UE as an ESM or DSM UE can be controlled from RADIUS via inclusion of *Alc-Wlan-Ue-Creation-Type* VSA. The VSA can be included in access-accept for a UE that is auto-signed-in (for example, it does not need web redirect to portal), or in a COA message triggered to remove web redirect for a UE after successful portal authentication. The VSA is described in the *7750 SR and VSR RADIUS Attributes Reference Guide*. If *Alc-Wlan-Ue-Creation-Type* is not present in access-accept (for auto-signed UE) or in the COA message (for UE creation of portal authenticated UE), then the UE is created as an ESM host. In this release DSM is not supported for UEs which require a GTP host. If *Alc-Wlan-Ue-Creation-Type* indicates a DSM UE then any IPv6 or GTP related parameters in access-accept or COA is ignored, and the UE is created as a DSM host. *Alc-Wlan-Ue-Creation-Type* cannot be changed mid-session via COA. A COA containing *Alc-Wlan-Ue-Creation-Type* for an existing UE does not result in any change of state, and is NACK'ed.

11.13.1 DHCP

Based on DHCP and L2 NAT configuration on the ISA, the configured IP address (**l2-aware-ip-address** configured under **vlan-tag-ranges range start** *vlan-id* **end** *vlan-id* or **vlan-tag-ranges range default**) is assigned to the user via DHCP. A different DHCP lease-time can be configured for an un-authenticated and an authenticated user for which an ESM or DSM host has been created. DHCP return options, for example, DNS and NBNS server addresses can be configured. This configuration can be per soft-wlan-gw group interface (by explicitly configuring it under **vlan-tag-ranges range default**), or per VLAN range (where a VLAN tag corresponds to an SSID). By default, for open SSIDs, DHCP DORA is completed, and authentication request is sent to AAA server only on reception of the first Layer 3 packet. However, with a **authenticate-on-dhcp** command configured under **vlan-tag-ranges range default** (default or specific range), authentication can be triggered on received DHCP DISCOVER or REQUEST when no UE state is present. If UE anchoring on GGSN/PGW is required, then **authenticate-on-dhcp must** be enabled, because the decision to setup GTP tunnel (in which case the IP@ for the UE comes from the GGSN/PGW) is based on RADIUS response.

In order to support unique inside IP addresses, the ISA Pool Manager can be used. Pools are allocated to each ISA in large blocks, requiring an IPv4 subnet with prefix-length 16. From this prefix, the subnet address (x.x.0.0), broadcast address (x.x.255.255), and gateway address (x.x.0.1) are reserved and not allocated to UEs, reducing the total number of available addresses by three. NAT pools are only available in non-retailer subscriber interfaces. Any retail service ID derived from configuration or AAA is only used for IPv6 pool selection and is ignored for IPv4 NAT pools. Forwarding in a different VRF can be achieved by selecting a different NAT policy and outside VRF.

11.13.2 Authentication and Accounting

The authentication is initiated from RADIUS client on the ISA anchoring the user, based on an isa-radius-policy (configured under **aaa**) and specified on the wlan-gw group-interface. This support exists in prior releases and is described in Authentication and Forwarding. The auth-policy can contain up to ten servers, five of which can be for authentication and all ten can be COA servers.

In order to generate accounting updates for DSM UEs, an accounting policy (type isa-radius-policy) must be configured under the **aaa** node and specified under **vlan-range** (**default** or **specific range**) on the wlan-gw interface. Accounting for DSM UEs includes **accounting-start**, **accounting-stop** and **interim-updates**. Interim-update interval is configurable under vlan-range on wlan-gw interface. The user-name format to be included in RADIUS messages is configurable in the auth-policy

and accounting-policy via the **user-name-format** command. By default, the user-name contains the UE MAC address, but can be configured to include the UEs MAC address and IP address, or circuit-id or DHCP vendor options. If **authenticate-on-dhcp** is enabled, then the IP address for the UE is not known prior to authentication, and, if the user-name is configured to contain both MAC and IP address, then only the MAC address is included.

The accounting-policy can be configured with attributes to be included in the accounting messages. The details of the attributes are covered in the *7750 SR-OS RADIUS Attributes Reference Guide*. The attributes are included here for reference.

```
*A:Dut-1>config>aaa# info
----------------------------------------------
  isa-radius-policy "isaRadiusPol1" create
      user-name-format mac mac-format alu
     acct-include-attributes
           acct-delay-time
           acct-trigger-reason
           called-station-id
           calling-station-id
           circuit-id
           dhcp-options
           dhcp-vendor-class-id
           frame-counters
           framed-ip-addr
           framed-ip-netmask
           hardware-timestamp
           inside-service-id
           mac-address
           multi-session-id
           nas-identifier
           nas-port-id
           nas-port-type
           octet-counters
           outside-ip
           outside-service-id
           port-range-block
           release-reason
           remote-id
           session-time
           subscriber-id
           ue-creation-type
           user-name
           wifi-rssi
           wifi-ssid-vlan
        exit
```

The **isa-radius-policy** for auth/COA and accounting specifies the server selection method for the servers specified in the policy with respect to load-balancing and failure of one or more servers. The three methods implemented include:

- Direct — Specifies that the first server is used as primary for all RADIUS messages, the second server is used as secondary (that is, used for all RADIUS messages if primary server fails), and so on.
- Round-Robin — RADIUS messages across accounting-sessions are distributed in a round-robin manner amongst the list of configured servers. All accounting messages for a given session are sent to the selected server for that session, until that server fails. If a server fails, then the sessions targeted to that server are distributed in a round-robin manner amongst the remaining servers. If the failed server comes back up, the sessions that were originally assigned to the failed server revert to the original server.
- Hash — Server is picked via hash on UE MAC. The hash list consists of all configured servers that are up. If a server fails, then the UEs hashed to that server are re-hashed over the remaining servers that are up.

If a response is not received for a RADIUS message from a particular server for a configurable timeout value (per server), and the time elapsed because the last packet received from this RADIUS server is longer than this configured timeout value, then the server is deemed to be down. Periodically an accounting-on message is sent to a server that is marked as down, to probe if it has become responsive. If a response is received then the server is marked as up.

```
*A:Dut-1>config>aaa# info
isa-radius-policy "isaRadiusPol1" create
            nas-ip-address-origin system-ip
            password "6mNsKxvTe.0.nNCTIpGFcu.rr/qtdijazQ3ED8WAFfk" hash2
            user-name-format mac mac-format alu release-reason
            servers
                access-algorithm hash-based
                retry 3
                router "Base"
                source-address-range 81.1.0.1
                timeout sec 5
                server 1 create
                    accounting port 1813
                    authentication port 1812
                    coa port 3799
                    ip-address 10.13.0.2
                    secret "3BmWbBfDO38hPY8DtLFn8bYDBaduy6w.ogeSUsouoHc" hash2
                    no shutdown
                exit
            exit
        exit
----------------------------------------------
*A:Dut-1>config>aaa#
```

11.13.2.1 DSM Data-Plane

NAT on the anchor ISA is required for forwarding of traffic to/from a DSM UE. There is no UE state in the IOM/IMM for a DSM UE. The downstream forwarding is based on FDB lookup that should match a route corresponding to the NAT outside pool, and get the downstream traffic to the right anchor ISA, where NAT is performed for the UE. The inside IP address assigned to the UE is the configured l2-aware-ip-address on the vlan-range (default or specific range) under wlan-gw interface. Therefore every UE corresponding to the default or specific vlan-range will get the same inside IP@. The NAT is L2-aware, and uses UE MAC to de-multiplex.

11.13.3 IP Filtering

Filtering based on protocol, destination IP, destination port or any combination is supported for traffic to and from the UE. The match entries and corresponding actions can be specified within the **isa-filter** which can be created in the **config>subscr-mgmt** context. The filter can be associated with a vlan-range (default or specific vlan-range) on wlan-gw interface, in which case all subscribers associated with the vlan-range is associated with an instance of this filter.

The supported filter actions include drop and forward. The first match will cause corresponding action to be executed and no further match entries is executed. In case there is no match or no action configured for a match, configurable default action for the filter is executed. The filter can be overridden on a per UE basis via RADIUS access-accept or COA. The new VSA *Alc-Wlan-Dsm-Ip-Filter* is defined for specifying the per UE filter from RADIUS. The VSA is defined in the RADIUS guide.

```
A:system>config>subscr-mgmt# info
----------------------------------------------
                isa-filter  "foo" type dsm create
                    default-action forward
                    entry 1 create
                        action drop
                        match protocol udp
                            dst-ip 239.0.113.0/32
                            dst-port  eq  53
                        exit
                    exit
                    ipv6
                        default-action forward
                        entry 1 create
                            action drop
                            match protocol tcp
                                dst-ip 2001:db8::/120
                                dst-port  eq  80
                            exit
                        exit
                    exit
```

```
-----------------------------------------------


*A:vsim>config>service>vprn# info
-----------------------------------------------
 subscriber-interface "s1" create
      group-interface "g1" wlangw create
             wlan-gw
                   vlan-tag-ranges
                          range default
                               distributed-sub-mgmt
                                    dsm-ip-filter “foo”

                               exit
                          exit
                    exit
               exit
-----------------------------------------------
```

11.13.3.1 Policing

Per UE policing for both ingress and egress direction is supported. Policers can be created in the **config>subscr-mgmt>isa-policer** context. The policers can be of type single-bucket (PIR) bandwidth limiting or dual-bucket (PIR and CIR) bandwidth limiting. Only policer action supported as permit-deny, non-conforming traffic is dropped, as opposed to marked out-of-profile. The administrative peak and committed rates and peak and committed burst sizes are configurable. For single-bucket bandwidth policers, CIR and CBS are not applicable, and only PIR and MBS are configurable.

```
*A:vsim>config>subscr-mgmt>isa-policer# info detail
-----------------------------------------------
                    no description
                    action permit-deny
                    cbs 100
                    mbs 200
                    rate 1000 cir 500
```

The policers can be associated with a vlan-range (default or specific vlan-range) on wlan-gw interface, in which case all subscribers associated with the vlan-range is associated with an instance of these policers. These ingress and egress policers can be overridden on per UE basis via RADIUS access-accept or COA. The new VSAs *Alc-Wlan-Dsm-Ingress-Policer* and *Alc-Wlan-Dsm-Egress-Policer* are defined for specifying the per UE policers from RADIUS. The VSAs are defined in the *7750 SR-OS RADIUS Attributes Reference Guide*. If the policers specified in access-accept are not found the message is dropped. If the policers specified in COA are not found, a NACK is sent back.

```
*A:vsim>config>service>vprn# info
----------------------------------------------
 subscriber-interface "s1" create
      group-interface "g1" wlangw create
             wlan-gw
                   vlan-tag-ranges
                          range default
                               distributed-sub-mgmt
                                    egress-policer "silver-egress"
                                    ingress-policer "silver-ingress"
                               exit
                          exit
                    exit
              exit
----------------------------------------------
```

11.13.3.2 Lawful Intercept (LI)

LI can be triggered for a DSM UE LI via CLI or RADIUS, and is performed post-NAT. Only routable encaps (IP/UDP/LI-shim) and IP-only mirror-dest are supported. A maximum of 2K DSM UEs per-chassis can be under LI simultaneously. LI mirror dest (service in which mirrored packets are injected) along with other required mirror information (mirror-dest type, encapsulation-type, ip-udp-shim, and encapsulation information, IP and UDP header information) is configurable. A DSM UE identified by its MAC address can be associated with the mirror destination (service in which mirrored packets for the host are injected) via the **li-source** command. For routable encapsulation (IP/UDP/LI-Shim), the session-id and transaction-id to be inserted in the LI-Shim are configured under **li-source**.

```
A:Dut-1>config>mirror# info
----------------------------------------------
        mirror-dest 60000 type ip-only create
            encap
                layer-3-encap ip-udp-shim create
                    gateway create
                        ip src 1.1.1.1 dest 2.2.2.2
                        udp src 2048 dest 2049
                    exit
                exit
            exit
            no shutdown
        exit

----------------------------------------------
A:Dut-1>config>li# info
----------------------------------------------
        li-source 60000
            wlan-gw
                dsm-subscriber mac 00:00:00:07:02:03
                    intercept-id 10000
                    session-id 20000
                exit
```

```
        exit
        no shutdown
    exit
```

LI can be enabled or disabled from RADIUS via inclusion of the *Alc-LI-Action* VSA in access-accept or COA. The *Alc-LI-Destination* VSA is required to indicate the mirror-dest service that the DSM UE under LI is associated with. The Intercept-Id and Session-Id for a DSM UE can be provided from RADIUS access-accept or COA via inclusion of Alc-LI-Intercept-Id and Alc-LI-Session-Id VSAs. These LI related VSAs are described in the *7750 SR and VSR RADIUS Attributes Reference Guide*.

Information for a particular li-source, and its associated mirror-dest can be shown via CLI.

11.13.3.3 Data-Triggered UE Creation

Similar to data-triggered UE creation with ESM, a DSM UE can also be created based on data-triggered authentication discussed in Data Triggered Subscriber Creation. The decision to create ESM versus DSM UE is based on the value of RADIUS VSA Alc-Wlan-Ue-Creation-Type present in the access-accept message. The data-triggered authentication and UE creation if configured provides for WLAN-GW IOM redundancy. The DSM UE is created on the standby ISA based on successful data-triggered authentication. Also, inter-chassis redundancy is supported for DSM UE based on data-triggered authentication, and is identical to ESM (as described in Data Triggered Subscriber Creation).

11.13.4 Idle-Timeout and Session-Timeout Management

The per UE idle-timeout value can be provided in RADIUS access-accept or COA for a DSM UE in standard Idle-Timeout attribute. The minimum idle-timeout allowed is 150 seconds. The idle-timeout is enforced on the ISA for a DSM UE. If there is no data to/from a UE for up to idle-timeout value, the UE is removed and accounting-stop is sent. Subsequently, if a UE re-associates and connects to an open SSID on an AP, and has an IP address with a valid lease, then the first data packet from the UE triggers authentication. Successful authentication results in creation of DSM UE.

To improve idle-timeout behavior an optional SHCV check can be performed after idle-timeout expires. This check verifies connectivity to all of the DHCP, DHCPv6 and /128 SLAAC addresses using ARP and/or NDP. While the check is performed for every address, the result is applied to the whole UE. The UE is only deleted when verification of all addresses fails. When at least one connectivity verification succeeds the UE and all of the allocated addresses are kept and the idle-timeout process is restarted.

The per UE session timeout value can be provided in RADIUS access-accept or COA in standard Session-Timeout attribute. The value is interpreted as an absolute value, and the UE is unconditionally deleted regardless of activity. The minimum allowed value for session-timeout is 300 seconds.

11.13.5 Operational Commands

The following shows the command usage to dump information on UE under LI (only allowed to users with LI privilege).

```
A:Dut-1# tools dump li wlan-gw ue
No sessions on Slot #2 MDA #1 match the query
===============================================================================
Matched 2 sessions on Slot #2 MDA #2
===============================================================================
UE-Mac          : 00:00:00:07:02:03     Mirror Service  : 60000
LI Intercept-Id : 10000                 LI Session-Id   : 20000
-------------------------------------------------------------------------------
UE-Mac          : 00:00:00:07:02:08     Mirror Service  : 60000
LI Intercept-Id : 42                    LI Session-Id   : 2013
-------------------------------------------------------------------------------
===============================================================================


A:Dut-1>show>li# li-source 60000
===============================================================================
Mirror Service
===============================================================================
Service Id       : 60000                Type          : ipOnly
-------------------------------------------------------------------------------
L3 encap type    : ip-udp-shim          Router        : Router: Base
                                        Direction bit : No
Primary gateway
Source IP        : 1.1.1.1              Dest IP       : 2.2.2.2
Source UDP port  : 2048                 Dest UDP port : 2049
===============================================================================
Local Sources
-------------------------------------------------------------------------------
Admin State      : Up
-------------------------------------------------------------------------------
WLAN Gateway LI sources
-------------------------------------------------------------------------------
MAC-Address                                          Intercept-Id Session-Id
```

```
-------------------------------------------------------------------------------
00:00:00:07:02:03                                    10000        20000
===============================================================================
```

11.13.6 Pool Manager

To support allocations of unique IP addresses each ISA is assigned pools from a centralized pool manager on the CPM. The ISA can subsequently assign addresses from these pools to UEs, but this state is not synchronized back to the CPM. Different applications have different pools, for example, SLAAC and DHCPv6 IA_NA cannot share a single pool. To support Wholesale/Retail scenarios a pool-manager can be configured per subscriber interface.

The allocation of additional pools and freeing up unused pools is based on configurable high and low watermarks. When the usage level of all pools combined on an ISA reaches the high watermark, a new pool is allocated. When the usage level of a single pool reaches zero and the usage level of the other pools combined is below the low watermark, this pool is freed.

With redundancy, the pool manager signals the pools that were allocated to the failed ISA back to the new active ISA. These pools can no longer be used to allocate new addresses because allocations are lost. However, these can still be used to forward traffic based on data-triggered UE creation. This is supported both for IOM redundancy and Active/Standby WLAN-GW redundancy. The new active ISA also receives new pools that it can use for new allocations.

The Pool Manager uses DHCPv6 Prefix Delegation to allocate pools to the ISAs. Each ISA is represented by a separate DHCPv6 Client ID. These clients request fixed prefix sizes to accommodate up to 64K UEs. With Active/Standby redundancy the Pool Manager uses a DHCPv6 Lease Query Message to retrieve the prefixes that were allocated to the failed WLAN-GW. In order to identify the correct PD leases in the DHCPv6 server, a configurable virtual-chassis-name is added to the DHCPv6 client-id, this value should be identical on both WLAN-GWs and unique otherwise. The Pool Manager will always send out a DHCPv6 Relay message and supports up to eight DHCPv6 servers.

IPv4 pools are supported by encoding an IPv4 subnet into an IPv6 prefix. The least significant 32 bits of the prefix are treated as an IPv4 address and the allocated prefix-length is subtracted with 96 to obtain an IPv4 prefix length. It is recommended to use the special IPv6 prefix “::ffff:/96” to provision these pools.

```
A:system>config>subscr-mgmt>wlan-gw# info
----------------------------------------------
            virtual-chassis-identifier "wlan_gw_pair"
----------------------------------------------
```

```
A:system>config>service>vprn>sub-if>wlan-gw# info
----------------------------------------------
                pool-manager
                    watermarks high 85 low 66
                    wlan-gw-group 1
                    dhcpv6-client
                        server 2001:db8::1
                        lease-query max-retry 2
                        slaac
                            pool-name "pool_ue_pd_v6_slaac"
                            no shutdown
                        exit
                        ia-na
                            pool-name "pool_ue_pd_v6_dhcp6"
                            no shutdown
                        exit
                        dhcpv4-nat
                            pool-name "pool_ue_pd_v6_dhcp4_nat"
                            no shutdown
                        exit
                    exit
                exit
----------------------------------------------
```

11.13.7 DHCPv6 and SLAAC

DHCPv6 and SLAAC support can be configured per VLAN range. Authentication can be triggered by a Router Solicit, DHCPv6 or DHCP packet but will only be triggered once per UE. Each UE can be assigned a unique DHCPv6 IA_NA address (/128) or SLAAC prefix (/64) from the ISA pools. The IPv6 pools are installed by the centralized pool manager. SLAAC privacy extensions are supported and up to three /128 SLAAC addresses can be learned via either Duplicate Address Detection or upstream data. Wholesale/retail is supported (IPv6 only) both via RADIUS and per vlan-range CLI, the applicable pool is selected from the retailer service. ESM and DSM IPv6 are not supported in the same *vlan-range* context.

```
A:system>config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range# info
----------------------------------------------
                                ...
                                authenticate-on-dhcp
                                ...
                                dhcp6
                                    active-preferred-lifetime hrs 1
                                    active-valid-lifetime hrs 1
                                    no shutdown
                                exit
                                slaac
                                    active-preferred-lifetime hrs 1
                                    active-valid-lifetime hrs 1
                                    no shutdown
                                exit
                                ...
```

```
-----------------------------------------------
```

Configuration of other DHCPv6/SLAAC parameters, such as server DUID and RA flags is taken from the **wlan-gw** *group-interface* configuration. For DSM only the configuration of the **wlan-gw** *group-interface* applies, the retailer interface cannot override this configuration.

```
A:system>config>service>vprn>sub-if>grp-if>ipv6# info
-----------------------------------------------
                        router-advertisements
                            other-stateful-configuration
                            prefix-options
                                autonomous
                            exit
                        exit
                        dhcp6
                            proxy-server
                                server-id duid-en string "example_duid"
                            exit
                        exit
```

A subset of DHCPv6 options retrieved by the pool-manager in the PD process is reflected in DHCPv6 towards the client. For IA_NA leases these are included in the associated DHCPv6 messages. For SLAAC allocations, the DNS option can be reflected in the Router Advertisement and all options can also be reflected in a stateless DHCPv6 Information Reply message.

When using a captive portal, different valid/preferred lifetimes can be configured for authenticated and un-authenticated UEs. The DHCPv6 lease time is equal to the applied valid-lifetime and can be extended via the regular renew process. SLAAC lifetime is equal to the applied valid-lifetime and is extended when sending an RA including the SLAAC prefix. To avoid infinite SLAAC allocations, when sending an unsolicited RA, SHCV is performed for all learned /128 addresses. If SHCV fails for all addresses, the unsolicited RA will not contain the SLAAC prefix and the SLAAC lifetime will not be extended.

In redundancy scenarios the new active ISA will migrate the leases in the old pools as soon as possible to a lease in the new pools. For SLAAC this is done by sending an unsolicited RA to deprecate the old prefix (lifetimes 0) and include a new prefix. For DHCPv6 this is done during the first Renew, that will again deprecate the old address (lifetimes 0) and include a new address in the same IA.

11.13.8 Application Assurance Support

To support Application Assurance (AA) with DSM, an AA group needs to be linked to a WLAN-GW group. By default, every ISA configured in the WLAN-GW, including standby ISAs, requires a counterpart in the AA group. In case AA is only required for a subset of UEs, an oversubscription factor can be configured. With an oversubscription factor of N, each AA ISA can serve up to N WLAN GW ISAs, including standby ISAs.

When oversubscription is used, the maximum number of AA-enabled UEs per WLAN-GW ISA is the maximum number of UEs for that ISA divided by the oversubscription factor.

AA is enabled per-UE by provisioning an application profile to the UE. A default profile can be provisioned under a **vlan-range**, or a specific profile can be signaled using RADIUS.

When an ISA AA fails, the corresponding BB ISAs continue to forward traffic without AA functionality. After the AA ISA recovers, AA is enabled again for those UEs.

For more information on how to configure AA, refer to the *7450 ESS, 7750 SR, and VSR Multiservice Integrated Service Adapter Guide*.

11.14 Call Trace

Call trace is supported for use on the ISA and the CPM. WLAN-GW-specific commands apply only to ISA tracing. To enable tracing on the CPM, IPoE session commands must be used without a SAP, circuit ID, or remote ID.

See Call Trace in the *Triple Play Enhanced Subscriber Management* chapter for more information.

11.15 Distributed RADIUS Proxy

The distributed RADIUS proxy acts just like the regular RADIUS proxy but runs on an ISA and is designed for high scale and high performance. It can handle a high number of RADIUS transactions, therefore it is able to keep up with EAP authentications that consists of many RADIUS transactions (EAP-PEAP) and all the accounting messages sent by an Access Point for a particular UE. The distributed RADIUS proxy is designed to handle the scale and performance of Distributed Subscriber Management (DSM) but can also be used as a performance improvement for Enhanced Subscriber Management (ESM). All common server-selection mechanisms are supported (direct, round-robin, hash-based) and both IPv4 and IPv6 RADIUS clients can communicate with the proxy. Important differences with the CPM based proxy are no IPv6 support towards the RADIUS server and no python support on any of the interfaces.

The distributed proxy also supports caching an access-accept to aid authentication of Layer 3 setup (DHCP/SLAAC/DHCPv6). After UE creation the cache supports tracking of both accounting and authentication messages. Contrary to the CPM-based RADIUS proxy the key used in the cache is always the calling-station-id attribute and it is expected to contain the UE MAC address, as specified in RFC 3580, *IEEE 802.1X Remote Authentication Dial In User Service (RADIUS) Usage Guidelines*. Accounting-on and accounting-off messages are not supported. The RADIUS proxy cache works with both ESM and DSM UEs.

For caching to work, the distributed proxy makes sure that all packets are routed via the anchor ISA tied to the UE. An AP will send a RADIUS packet to the radius-proxy IP address shared by all ISAs, the WLAN-GW will forward the packet to a distributor ISA based on the source IP address of the radius packet. That ISA then looks for the calling-station-id and forwards the packet to the correct anchor-isa to handle proxy functionality and caching. If no calling-station-id is found (such as acct-on/acct-off), the packet is always forwarded to a fixed ISA that is chosen at startup. The chosen ISA will forward the packet with a per-ISA IP as source-ip, this source-ip is assigned at startup from the range configured under **configure>aaa>isa-radius-policy** *policy-name*. From server to client the packet is sent back to that IP address and therefore immediately arrives at the correct anchor ISA, which subsequently forwards the packet straight to the AP without an additional ISA pass through.

Figure 185 **Distributed RADIUS Packet Forwarding**

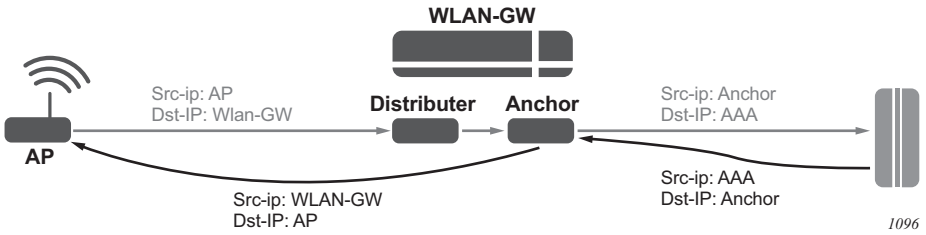

The following is a distributed proxy configuration example.

```
#--------------------------------------------------
/configure service vprn 50 radius-proxy
----------------------------------------------
    server "distributed_radius_proxy" purpose accounting authentication wlan-gw-
group 1 create
        cache
            key packet-type request attribute-type 31
            timeout min 5
            no track-accounting
            track-authentication accept
            track-delete-hold-time 0
            no shutdown
        exit
        default-accounting-server-policy "wlangw_isa_radius"
        default-authentication-server-policy "wlangw_isa_radius"
        no description
        no load-balance-key
        no python-policy
        secret "BLoAGDmsLt/Rs9LLU5/lESjjqZa/ssWnEIMJNvgBwmo" hash2
        send-accounting-response
        wlan-gw
            address 10.1.10.1
            ipv6-address 2001:db8::0
        exit
        no shutdown
    exit
----------------------------------------------
/configure aaa isa-radius-policy "wlangw_isa_radius"
----------------------------------------------
            password "rNPEv/V0j095N0Qy4rnekTVbF89OIlVj" hash2
            servers
                router "Base"
                source-address-range 10.100.100.4
                server 1 create
                    authentication
                    ip-address 10.100.100.2
                    secret "rNPEv/V0j095N0Qy4rnekPU0fmH2TwEl" hash2
                    no shutdown
                exit
            exit
----------------------------------------------
```

11.15.1 Enhanced Subscriber Management

For ESM support **authenticate-on-dhcp** must be enabled under **configure>service>vprn|ies** *svc*>**subscriber-interface** *sub-if*>**group-interface** *grp-if*>**wlan-gw vlan-tag-ranges range start** *start* **end** *end*. When receiving DHCPv4 the ISA will send the DHCP message and the cached access-accept to the CPM which will further process the setup sequence. On the CPM a regular radius authentication policy should be picked up for the UE either through configuration on the group-interface or via the LUDB. Typically this policy will reflect the ISA-policy. This policy is used as a context to store the access-accept on the CPM for 10s.

IPv6 hosts are supported but can only be authenticated after DHCPv4 has triggered the promote from ISA to CPM. When IPoE linking is enabled a SLAAC host is created together with the DHCPv4 host as usual. If an additional IPv6 host would arrive after the 10s timeout, a regular RADIUS authentication is started from the CPM using the previously mentioned RADIUS policy.

When tracking is enabled, the radius messages are handled on the ISA and specific tracking actions (mobility, delete) are sent directly to the CPM.

11.15.2 Distributed Subscriber Management

For DSM support the radius-proxy cache is directly tied to the UE record on the anchor ISA and is automatically used during UE creation. Tracking immediately executes the associated actions (mobility, timed host-delete) on the UE record. If a cached accept would time out before DHCP is received, a regular RADIUS authentication is used using the configuration under **configure>service>vprn|ies** *svc*>**subscriber-interface** *sub-if*>**group-interface** *grp-if*>**wlan-gw vlan-tag-ranges range start** *start* **end** *end*>**authentication**.

11.15.3 VLAN Awareness

The distributed RADIUS proxy optionally allows inclusion of the Alc-WLAN-SSID-VLAN attribute in an Access-Accept message. When this attribute is received, any subsequent traffic, including control plane traffic, must match the VLAN or VLAN-range, otherwise it is dropped. When an Access-Accept message with this attribute is received for a UE, one of the following scenarios may occur.

- The VLAN in authentication matches the current UE VLAN or VLAN-range. This is assumed to be a re-authentication within the same SSID and no further action is taken.

- The VLAN in authentication does not match the current VLAN or VLAN-range. It is assumed that the UE moved between SSIDs and that this is a new initial authentication. The UE is returned to an authenticated-only state from which the UE can be recreated with the correct new SSID parameters. If the UE was in the ESM or DSM state, it is gracefully removed. Any allocated IP addresses are released and final accounting data is sent.

The DRP checks on the exact VLAN or VLAN range based on configuration of inter-VLAN mobility. When inter-VLAN mobility is enabled, it is assumed that an SSID maps to a VLAN range instead of a single VLAN, and all comparisons are adapted accordingly. Additionally, when receiving an Access-Accept message for a UE with a different VLAN in the same VLAN range, the UE starts using the new VLAN to send downstream traffic.

This is recommended for deployment models where there are multiple SSIDs. Without specifying the VLAN, every authentication for a UE would be considered as a re-authentication for the same SSID and would make no changes to the WLAN-GW state. After authentication, the UE would send control-plane messages on the new SSID with a different VLAN or VLAN-range. The WLAN-GW would see this as a non-seamless change and trigger SHCV to remove the UE. Only after the UE is removed would control-plane traffic on the new VLAN or VLAN-range succeed, and it would also require an additional authentication trigger.

11.15.4 Operational Commands

The following commands will display all statistics related to the radius-proxy, both for communication towards the client and for communication towards the server.

show router *router-id* **radius-proxy-server** *server-name* **statistics**

clear router *router-id* **radius-proxy-server** *server-name* **statistics**

Example output:

```
*A:Dut-C# show router 50 radius-proxy-server "radius_proxy_isa" statistics
...
Group 1 member 3
-------------------------------------------------------------
Rx packet                                                  : 2
Rx Access-Request                                          : 2
Rx Accounting-Request                                      : 0
Rx dropped                                                 : 0
  Retransmit                                               : 0
  Wrong purpose                                            : 0
  No UE MAC to cache                                       : 0
  Client context limit reached                             : 0
  No ISA RADIUS policy configured                          : 0
```

```
  Invalid attribute encoding                          : 0
  Invalid password                                    : 0
  Accounting-Request with invalid Acct-Status-Type    : 0
  Accounting-Request with no Acct-Status-Type         : 0
  Invalid accounting Authenticator                    : 0
  Invalid Message-Authenticator                       : 0
  Management core overload                            : 0

Tx Access-Accept                                      : 1
Tx Access-Reject                                      : 0
Tx Access-Challenge                                   : 1
Tx Accounting-Response                                : 0
Tx dropped                                            : 0
  Server timeout                                      : 0
  Invalid response Authenticator                      : 0
  Invalid Message-Authenticator                       : 0
  Invalid attribute encoding                          : 0
  RADIUS server send failure                          : 0
...
```

The following RADIUS proxy messages sent to the server using this policy will also be counted here.

show aaa isa-radius-policy *policy-name*

clear aaa isa-radius-policy *policy*-**name statistics**

Example output:

```
*A:Dut-C# show aaa isa-radius-policy "wifi_isa_radius"
...
Server 1, group 1, member 3
-------------------------------------------------------------------------------
Purposes Up                                   : accounting authentication
Source IP address                             : 10.100.100.6
Acct Tx Requests                              : 0
Acct Tx Retries                               : 0
Acct Tx Timeouts                              : 0
Acct Rx Replies                               : 0
Auth Tx Requests                              : 2
Auth Tx Retries                               : 0
Auth Tx Timeouts                              : 0
Auth Rx Replies                               : 2
CoA Rx Requests                               : 0
...
```

11.16 WLAN-GW 1:1 Active-Backup Redundancy

This feature provides support for 1:1 inter WLAN-GW active-backup redundancy. The failure detection and switchover mechanism is contained in WLAN-GWs, and there is no dependency on the AP to detect failure of WLAN-GW and switch traffic to tunnel endpoint on a different WLAN-GW. There is also no dependency on NAT or a particular flavor of NAT on WLAN-GW. If local DHCP servers are used for address allocation, then DHCP leases in the server are synchronized to the backup WLAN-GW via MCS. However, ESM state for the UE is created on the backup WLAN-GW based on data-triggered authentication after switchover. The granularity of switchover is subscriber-interface. Both WLAN-GWs are required to be configured with the same tunnel endpoint address. Also, the subscriber interfaces on both WLAN-GW must be configured with the same subnets. Only the WLAN-GW that is deemed as active announces the tunnel endpoint address in routing towards the APs.

Active-backup decision is based on monitor and export route concept (same as what is used with NAT redundancy). Monitor and export routes are configured on the subscriber-interface on both WLAN-GWs. These should be complementary with respect to the ones on the other WLAN-GW. When WLAN-GW group goes up operationally, check is made in the FDB for presence of monitor route (which is the route exported by the other WLAN-GW). If it is not found, then the WLAN-GW assumes active state with respect to ownership of the tunnel end-point address, and the tunnel end-point address is announced in IGP towards the AP (subject to configured IGP and routing policy). The active WLAN-GW also announces the aggregate subscriber subnets upstream in routing. When WLAN-GW group comes up operationally, and detects the monitor route in the FDB, it assumes standby state with respect to the tunnel endpoint address. It does not announce the tunnel endpoint or the subscriber subnets in routing.

Each WLAN-GW will need to track the monitor route in the FDB. If the monitor route is no longer in the FDB, and the WLAN-GW is in standby state, it will transition to active, and announce the tunnel end-point towards APs, and subscriber subnets upstream. This will draw the traffic from the AP to the backup WLAN-GW. Redundancy is non-revertive. The monitor and export routes are configured on the subscriber-interface.

```
config>service>ies>sub-if
wlan-gw
redundancy
        [no] export <ip-prefix/length>
        [no] monitor <ip-prefix/length>
   exit
exit
```

If the number of operationally up WLAN-GW IOMs in wlan-gw group drops below the number of active IOMs configured, the WLAN-GW group is brought down (based on the **oper-down-on-group-degrade** command under the WLAN-GW interface), and switchover procedures for the subscriber-interface are triggered (export route, tunnel endpoint address and subscriber subnets are withdrawn from routing).

```
config>service>vprn>sub-if>grp-if
config>service>ies>sub-if>grp-if
wlan-gw
[no] oper-down-on-group-degrade
```

The switchover can also be triggered administratively on per subscriber-interface basis using the **tools perform** command.

```
*A:vsim-07-cpm# tools perform wlan-gw redundancy force-switchover service <service-
id> interface <ip-int-name>
```

11.16.1 DHCP Server Redundancy

1:1 redundancy provided with this feature only handles complete failure of WLAN-GW (either due to chassis reboot or due to number of operational WLAN-GW IOMs in WLAN-GW group falling below the number of active WLAN-GW IOMs, which will operationally bring down the WLAN-GW group, and trigger switchover). For any partial failures (port, MDA or IOM failure), it is assumed there is network level redundancy, such that the soft-GRE tunnel is re-routed to the primary WLAN-GW. This ensures there is only one active WLAN-GW owning the subnets defined on the two WLAN-GWs (that is, allows local/local subnets). The DHCP server(s) state is synchronized between the two WLAN-GWs using MCS.

Supported access includes:

- DHCPv4 Relay to external server
- DHCPv4 Relay to local server
  - Pool name could be returned by AAA (framed-pool) in access-accept
  - Pool name could come from LUDB (as relay we would set use-pool-from-client). LUDB could be specified under group-interface or under DHCP server. LUDB or AAA returned pool allows support for per SSID pool selection. SSID is contained in circuit-id.
  - Local pool selection based on giaddr
- DHCPv4 proxy (IP@ from AAA or IP@ from PGW/GGSN)

Unnumbered case should work both in relay and proxy scenarios. IPv6 is not supported in this release (as we don't support data-triggered auth and subscriber creation for IPv6). Therefore, DHCPv6 server synchronization is not applicable. Also, IPv4 address from LUDB is not supported in this release (as data-triggered authentication against LUDB is not supported).

11.16.2 Subscriber Creation after Switchover

When standby WLAN-GW transitions to active state, and receives data on the anchor ISA there will not be any UE state on the anchor ISA. Data-triggered authentication Data Triggered Subscriber Creation is used to create the subscriber. In order to infer how the UE originally obtained the IP@ (DHCP relay versus proxy, such as AAA or GTP), the following holds:

1. If any GTP related parameters are returned in access-accept, then it is assumed the IP@ comes from GGSN/PGW, and the origin for the IP@ is assumed to be GTP.
2. If no GTP parameter is returned, and access-accept contains framed-IP, then proxy case is assumed (that is, the origin as AAA).
3. If no GTP parameter or framed-IP is returned, then DHCP relay is assumed. The remaining lease time is set to initial lease-time (if it was originally provided from AAA on primary WLAN-GW, it could be provided in access-accept for data-triggered auth on backup WLAN-GW). If AAA does not provide it, then it is initialized to default value of 7 days.

If authentication indicates GTP for the subscriber, then create-session-request is signaled with Handover indication. However, for dual-sack subscriber over soft-GRE, if AAA returns the SLAAC prefix in access-accept (in response to IPv4 data-triggered auth), and linking is configured, RA message is sent (unicast to client's MAC@), and a SLAAC host is created.

11.17 WLAN-GW Triggered Stateless Redundancy (N:1)

Existing stateless redundancy, described in Data Triggered Subscriber Creation, is enhanced to support WLAN-GW based failure detection and switchover based on monitor and export route mechanism described above. The AP is not required to be configured with different tunnel endpoint addresses for active and standby WLAN-GWs. Single tunnel endpoint address is configured on the APs. The tunnel endpoint address is only announced in routing by the primary WLAN-GW as described in the section above. This form of redundancy as described in Data Triggered Subscriber Creation, required L2-aware NAT. After failure, the subscriber on the standby WLAN-GW that transitions to primary is based on data-triggered authentication. This is supported for both ESM and DSM.

11.18 AP Triggered Stateless WLAN-GW Redundancy (N:1)

Existing AP controlled redundancy, described in Data Triggered Subscriber Creation, is enhanced to trigger switchover on primary WLAN-GW if the number of WLAN-GW IOMs in the WLAN-GW group fall below number of active WLAN-GW IOMs. Based on a configuration **configure>service>vprn|ies** *svc*>**subscriber-interface** *sub-if*>**group-interface** *grp-if*>**wlan-gw** command, the WLAN-GW group is operationally brought down if a WLAN-GW IOM fails and the number of WLAN-GW IOMs fall below number of active WLAN-GW IOMs configured for the WLAN-GW group. This results in loss of route to the tunnel endpoint from the active WLAN-GW. The AP will detect this as WLAN-GW failure, and start tunneling the data to a configured backup WLAN-GW, where the subscriber is created based on data-triggered authentication. This is supported for both ESM and DSM.

11.19 IPv6-only Access

In order to accommodate IPv6 only AP/CPEs, IPv6 wlan-gw tunnel transport, and IPv6 client-side support for RADIUS-proxy have been added.

11.19.1 IPv6 GRE Tunnels

Support for IPv6 GRE tunnels require configuration of local IPv6 tunnel end-point address under wlan-gw configuration on the group-interface. The transport for L2oGRE (or L2VPNoGRE) packet is IPv6 as shown below. The outer IPv6 header contains the value 0x2F (GRE) in its Next Header field. GRE header contains protocol Ethernet (0x6558) or Ethernet-over-MPLS (0x8847) as in the case IPv4 GRE.

IPv6 Transport for L2oGRE Packet:

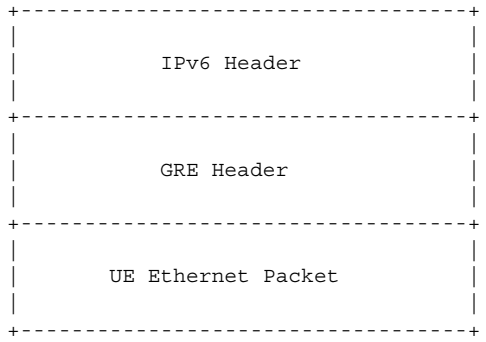

A single wlan-gw endpoint instance on the group-interface can have both IPv4 and IPv6 address configured as shown below. Inter-AP mobility between IPv4 and IPv6 only APs is supported in this scenario.

IPv6 Endpoint Configuration for WLAN-GW:

```
service
    vprn 300 customer 1 create
        group-interface "grp-intf-1" wlangw create
            wlan-gw
                gw-address 10.1.1.4
                gw-ipv6-address 2001:db8::0
                mobility
                   hold-time 0
                   trigger data iapp
                exit
                egress
                   shaping per-tunnel
```

```
                exit
                tcp-mss-adjust 1000
                vlan-tag-ranges
                    range start 100 end 100
                        data-triggered-ue-creation
                        retail-svc-id 402
                    exit
                exit
                router 30
                wlan-gw-group 1
                no shutdown
            exit
        exit
    exit
exit
```

The data-path for IPv6 GRE tunneled packets, including load-balancing of tunneled packets amongst set of ISAs in the WLAN-GW group, and anchoring after tunnel decapsulation remains unchanged. Per tunnel traffic shaping is supported similar to IPv4 tunnels. All existing per tunnel configuration on the group-interface described in previous sections (including mobility, egress shaping, VLAN ranges, and so on) is supported identically for IPv6 tunnels. Tunnel reassembly for upstream tunneled traffic is not supported for IPv6 tunnels in this release. TCP mss-adjust is supported for IPv6 tunnels, and is configurable under wlan-gw mode on group-interface. APs must use globally routable addresses for GRE IPv6 transport. Packets with extension headers are dropped.

11.19.2 IPv6 Client-Side RADIUS Proxy

RADIUS proxy is extended to listen for incoming IPv6 RADIUS messages from IPv6 RADIUS clients on AP/CPEs. The listening interface that the RADIUS proxy binds to must be configured with an IPv6 address as shown below. The IPv6 RADIUS proxy is solely for DHCPv4-based UEs behind IPv6 only AP/CPEs (IPv6-capable UEs are not supported in this release). All RADIUS-proxy functions (including caching, correlation with DHCPv4, and mobility tracking) are supported identically to existing IPv4 client-side RADIUS-proxy.

Configuration for IPv6 Client-Side RADIUS Proxy:

```
service
    vprn 300 customer 1 create
        shutdown
        interface "listening_radius_server" create
             address 10.9.9.9/32
             ipv6
                 address 2001:db8::0
             exit
             loopback
         exit
```

```
radius-proxy
    server "radius-proxy" purpose accounting authentication create
        shutdown
        cache
            key packet-type request attribute-type 31
            track-accounting stop interim-update accounting-on accounting-off
            no shutdown
        exit
        default-accounting-server-policy "radius_server_policy"
        default-authentication-server-policy "radius_server_policy"
        interface "listening_radius_server"
        load-balance-key attribute-type 102 vendor 5
        secret "AQepKzndDzjRI5g38L3LbbN3E8qualtn" hash2
        send-accounting-response
        no shutdown
    exit
exit
```

11.19.3 Dual-Stack UEs over WLAN-GW

This feature adds support for dual-stack UEs over wlan-gw. Each dual-stack UE appears to WLAN-GW as a bridged client. Dual-stack UE support includes both SLAAC and DHCPv6, with and without linking with DHCPv4. Handling of DHCPv6, RS/RA, and NS/NA messages over wlan-gw has been added. WLAN-GW can assign /128 GUA to the UE via DHCPv6 and/or assign a /64 prefix in SLAAC to each UE. Each UE can be handed via DHCPv6, a /128 IA_NA from a unique /64 prefix, with the “on-link” flag is off in the RA message. This is because the public WIFI users are distinct subscribers, and the communication must always be via WLAN-GW. The CPE must prevent local-switching on the WIFI link even if the /64 prefix is signaled as on-link or if the UEs are handed out /128 from the same /64 prefix.

Existing ESMv6 support on normal group-interface is applicable to wlan-gw group-interface, and is already documented in general ESMv6 sections in this guide. There are a few exceptions that are mentioned in sections below.

11.19.3.1 SLAAC Prefix Assignment

SLAAC prefix assignment to the UE can be from local prefix pool, where pool name can come from RADIUS in Alc-SLAAC-IPv6-Pool VSA or from LUDB (see general section on ESMv6 SLAAC pool assignment). Alternatively, the SLAAC prefix can be provided from RADIUS (in standard Framed-IPv6-Prefix attribute) or from LUDB. SLAAC with stateless DHCPv6 (DHCPv6 information-request) is supported. DNS can be sent in RA messages (per RFC 6106). RS authentication (based on MAC

address) can be configured (as described in general ESMv6 section on SLAAC only ESM hosts). SLAAC host is created on successful RS authentication. For successfully authenticated SLAAC host, an RA is sent in response to every received RS message (subject to a configured min-auth-interval). RA messages are sent to unicast MAC address of the UE.

SLAAC host creation can be linked to DHCPv4 by configuring **ipoe-linking** under group-interface. With **ipoe-linking** enabled, any received RS messages are dropped till DHCPv4 successfully authenticates and ESMv4 host is created. If **gratuitous-rtr-adv** is configured under ipoe-linking context then an RA is sent when ESMv4 host is created. If available, the SLAAC-prefix is included in the RA message. **shared-circuit-id** command under wlan-gw is not supported on wlan-gw interfaces. The O-Bit (other-stateful-configuration) is configurable on the group-interface.

11.19.3.2 DHCPv6 IA_NA Assignment

If UE requests DHCPv6 IA_NA, a /128 address can be provided from a unique /64 prefix per UE from a local-pool. The pool name can be provided from LUDB or from RADIUS (in Framed-IPv6-Pool attribute). The address could also be provided via LUDB or RADIUS (in Alc-IPv6-Address VSA). DHCPv6 can also be linked with DHCPv4 by enabling **ipoe-linking** command. The M-bit in RA message is configurable. DHCPv6 IA_NA is allowed if it is received after a SLAAC host exists, if **allow-multiple-wan-addresses** is enabled under group-interface ipv6 configuration. In previous releases, this is precluded. This however consumes two hosts (one each for IA_NA and SLAAC) per UE. Based on a configuration command **override-slaac**, SLAAC host can be deleted if DHCPv6 IA_NA host is successfully created. Prefix-delegation is not supported with DHCPv6 on wlan-gw interfaces.

11.19.3.3 Migrant User Support

Migrant user support is only applicable to IPv4. However, if linking is configured for SLAAC or DHCPv6 with DHCPv4 then RS and DHCPv6 messages are dropped till IPv4 ESM host exists (that is, the UE is out of migrant state). Once the IPv6 ESM host exists, that is, UE is out of migrant state, RA is sent to the UE (unicast MAC), and subsequent RS or DHCPv6 messages can result in creation of IPv6 ESM host. Therefore, with migrant UEs, linking should be enabled. SLA-profile instance accounting (with interim-updates), and per-host accounting (w/ interim-updates) are supported.

11.19.3.4 Accounting

Per SLA-profile instance accounting (with interim-updates) and per SLA-profile instance accounting (with interim-updates) with host accounting enabled is supported. The interim-updates are scheduled updates, and carry IPv4 address and IPv6 address or prefix are assigned to the UE.

A sample sequence with per SLA-profile instance accounting (with interim-updates) is shown below:

0. IPv4oE host created based on DHCPv4.

1. Acct-start generated (contains framed-ip-address).

2. SLAAC host comes up.

3. Next scheduled interim-update (contains framed-ip-address and framed-IPv6-Prefix, that is, SLAAC-prefix).

4. DHCPv6 IA_NA gets assigned and corresponding host is created.

5. Next Scheduled interim-update (contains framed-ip-address, framed-IPv6-Prefix and Alc-Ipv6-Address).

6. SLAAC host times out.

7. Next Scheduled interim-update (contains only Alc-IPv6-Address and will not contain framed-IPv6-Prefix).

8. DHCPv6 IA_NA lease times out.

9. Next Scheduled interim-update (contains only framed-ip-address).

A sample sequence with per SLA-profile instance accounting (with interim-updates) with host accounting enabled is shown below:

0. IPv4oE host created based on DHCPv4.

1. Acct-start for sla-profile instance generated (contains framed-ip-address).

2. Acct-start for DHCPv4 host is generated (contains framed-ip-address).

3. SLAAC host comes up.

4. Acct-start for SLAAC host is generated (this should contain framed-IPv6-Prefix, that is, SLAAC-prefix)

5. Next scheduled interim-update for sla-profile instance accounting (contains framed-ip-address and framed-IPv6-Prefix, that is, SLAAC-prefix).

6. DHCPv6 IA_NA gets assigned and corresponding host is created.

7. Acct-start for DHCPv6 IA_NA host is generated (contains Alc-Ipv6-Address).

8. Next Scheduled interim-update (contains framed-ip-address, framed-IPv6-Prefix and Alc-Ipv6-Address).

9. SLAAC host times out.

10. Acct-stop (SLAAC-host-acct-session-id) is generated.

11. Next Scheduled interim-update for sla-profile instance accounting (contains only Alc-IPv6-Address).

12. DHCPv6 IA_NA lease times out.

13. Acct-stop (DHCPv6-IA_NA-host-acct-session-id) is generated.

14. Next Scheduled interim-update for sla-profile instance accounting (contains framed-ip-address).

15. DHCPv4 lease times out.

16. Acct-stop (DHCPv4-host-acct-session-id) is generated.

17. Acct-stop for sla-profile instance accounting is generated.

11.20 Layer 2 Wholesale

This feature adds support for mapping a UE to a VPLS instance based on configuration. The mapping is explicitly created by assigning a VPLS instance to an SSID that the UE is connected to. The SSID is represented by the .1q tag in the received Layer 2 frames from the UE. A VPLS instance is configured per vlan-range on wlan-gw group-interface (as shown in Standalone WLAN-GW). This feature therefore enables Layer 2 wholesale, where traffic from all UEs on a particular SSID is transparently forwarded into the corresponding VPLS instance associated with the retail ISP. UE authentication, address assignment, Layer 3 classification and QOS are managed by the retail provider terminating the subscriber. There is no local-switching on the WLAN-GW providing the wholesale service. When a VPLS instance is configured under a VLAN-range, an internal SAP is implicitly created in the VPLS instance between each ISA and corresponding carrier IOM in the WLAN-GW group. The internal SAP is associated with an implicitly created SHG to constrain broadcast and multicast traffic received from UEs, such that it is not forwarded back on the SAP. Layer 2 wholesale and Layer 3 termination are possible simultaneously on same wlan-gw interface, because Layer 2 wholesale or Layer 3 termination is a per SSID decision. UE state on the ISA is removed when the UE MAC in the VPLS instance ages out based on **local-age** configured under VPLS service.

A **vpls-sap-template** command (described in the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN*) can be defined under **config>service>template** and associated with the VPLS service for Layer 2 wholesale via **config>service>vpls>wlan-gw>sap-template** command. Ingress and egress filter and QoS specified in the **vpls-sap-template** for the VPLS service is applied to the implicitly created internal SAP (between ISA and carrier IOM) in the VPLS service.

```
*A:vsim>config>service>vprn# info
----------------------------------------------
 subscriber-interface "s1" create
       group-interface "g1" wlangw create
              wlan-gw
                    vlan-tag-ranges
                           range start 100 end 100
                                l2-service 600
                                      no shutdown
                                exit
                           exit
                    exit
               exit
----------------------------------------------

*A:vsim>config>service>vpls# info
----------------------------------------------
 wlan-gw
     shutdown
     sap-template "foo"
```

```
 exit
-----------------------------------------------
```

11.21 VLAN to WLAN-GW IOM/IMM Steering via Internal Epipe

This feature provides the steering of traffic received on an access VLAN or spoke SDP from a WIFI AP/AC to a WLAN-GW IOM/IMM via an internal Epipe. The benefit of this internal steering is that all existing features available with native soft GRE tunnels on WLAN-GW IOM/IMM are now available to pure Layer 2 access via VLANs or spoke SDPs. The access SAP can be null, .1q, or q-in-q. Access SAPs aggregating WIFI APs or ACs can and be configured in the **configure>service>ies>sub-if>grp-if>wlan-gw>l2-access-points>l2-ap** or **configure>service>vprn>sub-if>grp-if>wlan-gw>l2-access-points>l2-ap** context

The aggregation network can insert up to two AP identifying VLAN tags, and the AP can insert a .1q tag (typically for identifying the SSID). The number of AP identifying tags sent on the internal Epipe depends on the encapsulation on the access SAP. For example, if an aggregation network inserts two AP identifying tags, and an access SAP is configured with null encapsulations, then the traffic sent on the internal Epipe will carry two AP identifying tags. The number of AP identifying tags in the frame forwarded over the internal Epipe must be configured via the **l2-ap-encap-type** command.

```
configure service (vprn|ies) <svc-id> subscriber-interface <sub-itf> group-
interface <grp-itf> wlan-gw
    l2-access-points
        [no] l2-ap <sap-id> [create]
            [no] encap-type {default|null|dot1q|qinq}
            [no] epipe-sap-template <name>
            [no] shutdown
        exit
    exit
    [no] l2-ap-encap-type {null|dot1q|qinq}
exit
```

The traffic on an internal Epipe is load-balanced among ISAs in the WLAN-GW group. The load balancing uses a hash based on AP identifying tags that remain on the frame after being received on the access SAP (based on the SAP encapsulation). This ensures all traffic from a particular AP is Epiped to the same ISA. Ingress and egress QoS and filters can be defined in an **epipe-sap-template** as displayed in the

configuration shown below, and associated with the access sap or spoke SDP. IP filters and DSCP remarking are not supported if more than two tags are present in the frame ingressing the SAP. Also, downstream filters and DSCP remarking is not applied if a retail tag is present. Both Layer 3 ESM and DSM as well as Layer 2 wholesale are supported for steered traffic.

```
configure service template epipe-sap-template <name> [create]
    egress
        [no] filter
            [no] ip <filter-id>
            [no] ipv6 <filter-id>
            [no] mac <filter-id>
        exit
        [no] qos <policy-id>
    exit
    ingress
        [no] filter
            [no] ip <filter-id>
            [no] ipv6 <filter-id>
            [no] mac <filter-id>
        exit
        [no] qos <policy-id> {shared-queuing|multipoint-shared}
    exit
exit
```

Currently, mobility from an AP that is reached over a VLAN or spoke SDP to an AP that is reached over a soft GRE or soft L2TPv3 tunnels are not supported. Each internal Epipe takes away two SAPs on each WLAN-GW IOM (one per ISA) in WLAN-GW group. With 64K SAPs per IOM, the maximum number of internal Epipes supported per chassis is 32K.

11.22 Soft-L2TPv3 Tunnels

This feature adds support for Layer 2 over soft-L2TPv3 tunnels. L2TPv3 is over UDP and both IPv4 and IPv6 transport is supported. The encapsulation with UDP allows NAT traversal. Soft-L2TPv3 tunnels are terminated on WLAN-GW IOM/IMM. All features supported with soft-GRE tunnels are supported identically with soft-L2TPv3 tunnels. L2TPv3 tunnels are stateless and there is no support for control channel, dynamic exchange of session-id and cookie, and L2-specific sub-layer for sequencing. Received cookie in L2TPv3 is reflected back. The AP can encode it's MAC address in 8-byte cookie. Based on configuration, the cookie can be ignored and just reflected back, or parsed to interpret AP-MAC from the least significant 6 bytes. Both L2TPv3 over IP and L2TPv3 over UDP encapsulation is supported. L2TPv3 tunnels are load-balanced from ingress IOMs to WLAN-GW IOMs based on source IP address. Figure 186 and Figure 187 shows these encapsulations with IPv6.

Figure 186 **L2TPv3 over UDP (IPv6 Transport)**

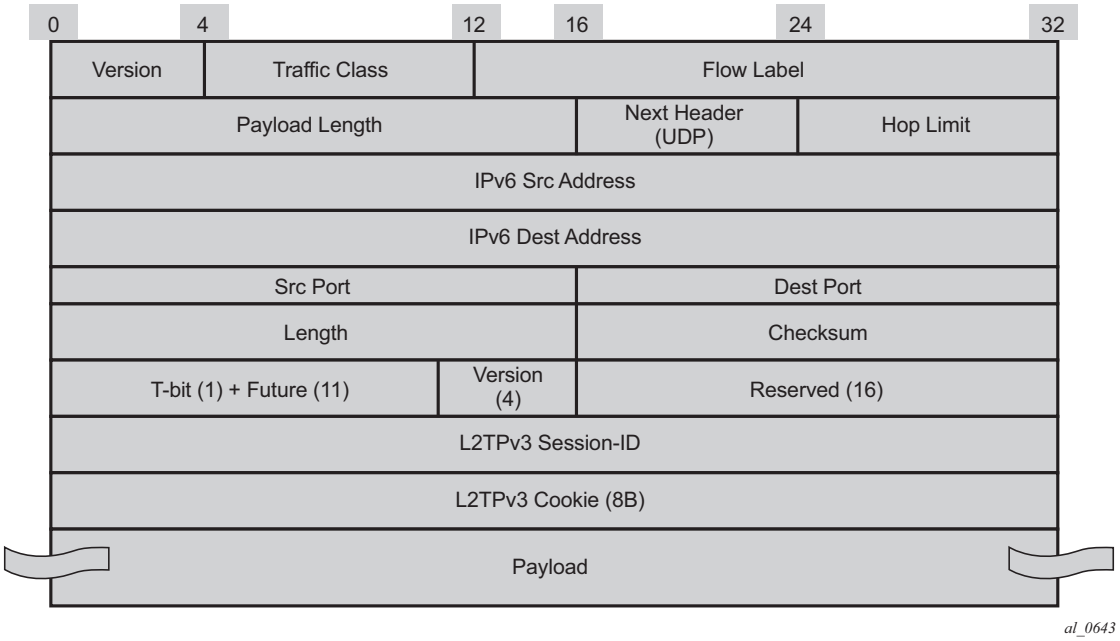

Figure 187 **L2TPv3 over IP (IPv6 Transport)**

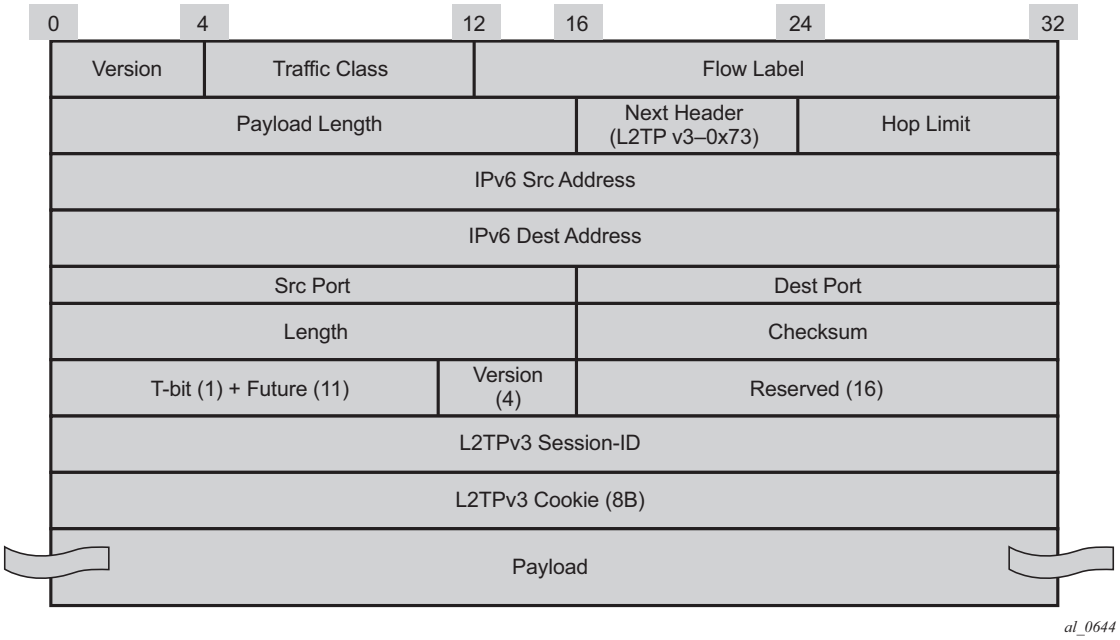

Enabling **multi-tunnel-type** on a **wlan-gw group-interface** allows multiple tunnel types (such as soft-GRE and L2TPv3) to the same gateway tunnel endpoint. Mobility between APs reachable with soft-L2TPv3 tunnels and APs reachable by soft-GRE tunnels is supported. There is feature and scale parity between soft-GRE and soft-L2TPv3 tunnels. The local tunnel gateway endpoint and other configurations parameters are shown below.

```
A:Dut-C>config>service>vprn>sub-if>grp-if>wlan-gw# info
----------------------------------------------
                    gw-address 10.1.1.3
                    gw-ipv6-address 2001:db8::0
                    learn-ap-mac
                    mobility
                        hold-time 0
                        multi-tunnel-type
                        trigger data iapp
                    exit
                    tunnel-encaps
                        learn-l2tp-cookie always
                    exit
                    wlan-gw-group 1
                    no shutdown
----------------------------------------------
```

11.23 WLAN-GW — Dynamic Tunnel X-Connect for Seamless Inter-WLAN-GW Mobility

This feature adds support for seamless inter WLAN-GW mobility when a UE roams from an AP behind a WLAN-GW to an AP behind another WLAN-GW. Prior to this feature, inter-WLAN-GW mobility was supported by creating a UE state on the target WLAN-GW based on a data-trigger, and carries over its L2-aware NAT inside IP address to the target WLAN-GW. However, the outside NAT pool (and hence the outside IP) changes, and any traffic for existing sessions is dropped. This feature introduces the concept of “home WLAN-GW” (H-GW) that provides an anchor point for a UE when it roams to an AP behind a different WLAN-GW, referred to as “visited WLAN-GW” (V-GW). With anchoring on the H-GW, the UE retails both its NAT inside and outside IP addresses. Existing NAT flows on the H-GW stay intact, and forwarding is seamless. The V-GW tunnels the UE traffic to the H-GW that anchors the UE.

Figure 188 **Dynamic Tunnel X-Connect for Seamless Inter-WLAN-GW Mobility**

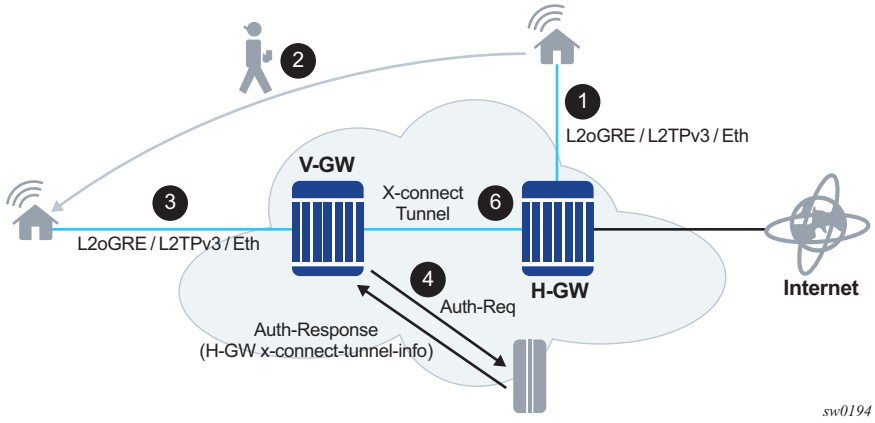

11.23.1 Processing on the V-GW

A control or data packet is received on the V-GW when a UE moves to an AP behind the V-GW. The packet can be received on any supported access type (L2oGRE, L2TPv3 tunnel, or internal VLAN anchor). The packet, after tunnel decapsulation, is forwarded to the anchor ISA on the V-GW. Authentication is triggered from the anchor ISA based on the configured ISA authentication policy. The **authenticate-on-dhcp** command must be enabled on V-GW for DHCP-triggered UE state creation scenarios to work.

If the access-accept message contains any VSAs related to x-connect tunneling, then the UE is created in the x-connect mode, and any subsequent Layer 2 frames received in the V-GW from the UE are transparently tunneled to H-GW. By default, the x-connect tunnel used between the V-GW and the H-GW is of same type (LoGRE or L2TPv3) as an access tunnel (for example, the tunnel from the AP to the V-GW). If the access type is L2-AP (such as an internal VLAN anchor), then the x-connect tunnel used is L2oGRE. An x-connect tunnel type can be overridden from the AAA in an access-accept message by a VSA (Alc-Xconnect-Tunnel-Type). IPv6 transport is used for x-connect tunnels.

The tunnel destination for an IPv6 address on the H-GW is returned in a VSA in an access-accept message (Alc-Xconnect-Tunnel-Ipv6). The service in which the tunnel is routed is also returned in a VSA (Alc-Xconnect-Tunnel-Service). The anchor ISA on the V-GW allocates a unique IPv6 address per UE for the x-connect tunnel source. This address is allocated from a contiguous IPv6 address range assigned to the ISA (based on the configured /64 prefix for the WLAN-GW group supporting x-connect tunneling). If the UE moves from one AP to another on the V-GW, then the anchor ISA allocates a new IPv6 address to use as tunnel sources for the x-connect tunnel. This change in x-connect tunnel source services as a UE mobility indication on the H-GW.

11.23.2 Processing on H-GW

Normal UE processing for WLAN-GW is performed on the H-GW. After tunnel decapsulation, the UE is received on the anchor ISA. If the UE state already exists on the anchor ISA, then it is treated as a mobility event, and the tunnel information in the UE (to the tunnel toward the V-GW). A mobility-triggered interim-update, if configured, is generated. The AP's MAC address must be reflected in **called-station-ID** attribute. The AP MAC is learned from the DHCP circuit-ID. However, if the mobility trigger is a data packet and the tunnel type is L2oGRE, then an ARPoGRE is generated (if configured) to learn the AP MAC. The V-GW forwards

ARPoGRE is received on x-connect tunnel to the access tunnel towards the AP, and reflects the received response back on x-connect tunnel. If the x-connect tunnel is type L2TPv3, then the AP MAC address can be carried in the cookie field. If the UE state does not exist, then normal UE processing leads to authentication and subsequent UE creation (as DSM or ESM UE).

IPv6 tunnel reassembly is not supported.

11.23.3 Idle Timeout Handling

An idle timeout (with optional SHCV) is independently enforced on the H-GW and V-GW. The idle timeout value can be provided by AAA. The UE state is deleted if the idle timeout triggers and there is no response to SHCV.

11.23.4 Distributed RADIUS Proxy for Closed SSID

An existing distributed RADIUS proxy (as described in Distributed RADIUS Proxy) can be enabled on both the H-GW and V-GW. Authentication and accounting messages are proxied to configured AAA servers. When the UE state exists on the V-GW, any authentication and accounting messages are visible on the H-GW, and are proxied directly to the configured AAA server from the V-GW. Data-triggered mobility must be enabled on the H-GW, as it can not rely on authentication and accounting messages to infer UE mobility when the UE is on V-GW.

11.23.5 H-GW Redundancy

Existing WLAN-GW controlled redundancy (mastership decision based on monitor or export route tracking mechanism, as described in section WLAN-GW 1:1 Active-Backup Redundancy) is supported. The backup H-GW advertises the tunnel endpoint IP in routing when it becomes the master to draw the traffic from the V-GW. The UE state is created on the new master based on data-triggered authentication.

11.24 ISA Operational Commands and Key Performance Indicators

This section provides an overview of operational commands that can be used to monitor ISA behavior and performance. This focuses on WLAN-GW functionality. If NAT is also used on the same ISA, refer to the NAT guides for similar commands.

11.24.1 ISA Resources

Resources on ISA BB are statically allocated at boot time and therefore strict have limits. These limits depend on both the provisioned ISA type used and on specific configuration flags. Use the command **tools dump wlan-gw isa resources mda** *mda-id* to generate a list of all resources, together with the currently in-use, free, maximum, and peak usage counters.

Most resources are directly related to configuration, and the corresponding configuration will fail if no resources are available. However, specific resources, such as the number of UEs on an ISA, are allocated dynamically. To facilitate operational maintenance, configure watermarks on those resources using the command **configure isa wlan-gw-group** *wlan-gw-group-id* **watermarks**. When the corresponding high watermark is reached, a notification is generated, which can prompt an operator to investigate the high resource consumption and take appropriate steps, such as optimizing load or adding additional resources.

11.24.2 ISA Load

Use the command **tools dump wlan-gw isa performance mda** *mda-id* **last** *time-span time-unit* to generate a high-level overview of ISA processing load and total data processed in a specified time span with a polling granularity depending on that time span. The following time spans and granularity are supported:

- last minute with second granularity
- last hour with minute granularity
- last day with hour granularity
- last month with day granularity

11.24.3 Packet Statistics

The ISA BB keeps extensive counters for each packet that is sent, terminated, or forwarded. These statistics are displayed using the command **show isa wlan-gw-group** *wlan-gw-group-id* **member** *member-id* **statistics**. Most statistics are straight-forward operational counters, such number of DHCP messages received or sent; however, there is also an extensive set of error counters available. When a WLAN-GW shows unexpected behavior, these counters can provide preliminary information about the issue. During debugging, it is recommended to first reset statistics in order to reset old and potentially unrelated error counters.

If a specified statistic displays an unexpected increase, it is not always easy to identify which UE or which behavior in a UE caused this increase. In this case, enable debugging of a specific statistic using the command **debug wlan-gw group** *wlan-gw-group-id* **statistics**. This command generates a hex dump of the first packet causing the statistic to increase in the debug logs. This packet can then be further investigated to determine the root cause and can be used to derive a UE MAC to start UE-based debugging.

11.25 WiFi Command Reference

11.25.1 Command Hierarchies

- WLAN-GW Commands
- ISA Commands
- WLAN-GW Service Commands
- Data Plane Related Commands
- RADIUS Server and Proxy Commands
- LUDB Matching for RADIUS Proxy Cache
- Port Policy Commands
- WIFI Aggregation and Offload – Migrant User Support Commands
- Show Commands
- Debug Commands
- Tools Commands
- Clear Commands

11.25.1.1 WLAN-GW Commands

Note: The **wlan-gw** commands apply only to the 7750 SR platform.

```
config
    — subscriber-mgmt
        — wlan-gw
            — distributed-sub-mgmt
                — gtp-peer-clear-timeout seconds
                — no gtp-peer-clear-timeout
            — virtual-chassis-identifier dual-homing-key
            — no virtual-chassis-identifier

configure
    — router
        — wlan-gw
            — distributed-sub-mgmt
                — ipv6-tcp-mss-adjust segment-size
                — no ipv6-tcp-mss-adjust
            — mobility-triggered-acct
```

```
— interim-update include-counters [hold-down seconds]
— interim-update
— no interim-update
```

11.25.1.2 ISA Commands

```
config
    — subscriber-mgmt
        — isa-filtername [type {dsm}] [create]
        — no isa-filter name
            — default-action {drop | forward}
            — no default-action
            — description description-string
            — no description
            — entry entry-id [create]
            — no entry entry-id
                — action {drop | forward | none}
                — no action
                — description description-string
                — no description
                — match protocol {any | icmp | tcp | udp | gre}
                — no match
                    — dst-ip ip-prefix/length
                    — no dst-ip
                    — dst-port operator port-number
                    — no dst-port
                    — src-ip ip-prefix/length
                    — no src-ip
                    — src-port operator port-number
                    — no src-port
                — ipv6
                    — default-action {drop | forward}
                    — no default-action
                        — entry entry-id [create]
                        — no entry entry-id
                            — action {drop | forward | none}
                            — no action
                            — description description-string
                            — no description
                                — match protocol {any | icmp | tcp | udp}
                                — no match
                                    — dst-ip ip-prefix/length
                                    — no dst-ip
                                    — dst-port operator port-number
                                    — no dst-port
                                    — src-ip ip-prefix/length
                                    — no src-ip
                                    — src-port operator port-number
                                    — no src-port
        — isa-policer policer-name [type policer-type] [create]
        — no isa-policer policer-name
```

- **action** {**permit-deny** | **priority-mark**}
- **no action**
- **adaptation-rule pir** {**max** | **min** | **closest**} [**cir** {**max** | **min** | **closest**}]
- **no adaptation-rule**
- **cbs** *burst-size*
- **no cbs**
- **description** *description-string*
- **no description**
- **mbs** *burst-size*
- **no mbs**
- **rate** *rate* [**cir** *rate*]
- **no rate**

config service

- **ies** *service-id*/**vprn** *service-id*
- **subscriber-interface** *ip-int-name*
  - **group-interface** *ip-int-name* [**create**]
  - **group-interface** *ip-int-name* [**create**] **lns**
  - **group-interface** *ip-int-name* [**create**] **wlangw**
  - **no group-interface** *ip-int-name*
    - **wlan-gw**
      - **vlan-tag-ranges**
        - **range start** [*range*] **end** [*range*]
        - **range default**
        - **no range start** [*range*] **end** [**range**]
        - [**no**] **authenticate-on-dhcp**
          - **authentication**
            - **authentication-policy** *policy-name*
            - **no authentication-policy**
            - **hold-time** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
            - **no hold-time**
            - **vlan-mismatch-timeout** *seconds*
            - **no vlan-mismatch-timeout**
          - **distributed-sub-mgmt**
            - **accounting-policy** *policy-name*
            - **no accounting-policy**
            - **accounting-update-interval**
            - **accounting-update-interval** [*interval*]
            - **no accounting-update-interval**
            - **def-app-profile**
            - **def-app-profile** *profile-name*
            - **no def-app-profile**
            - **dsm-ip-filter**
            - **dsm-ip-filter** *dsm-ip-filter-name*
            - **no dsm-ip-filter**
            - **egress-policer**
            - **egress-policer** [*policer-name*]
            - **no egress-policer**
            - **ingress-policer**
            - **ingress-policer** *policer-name*
            - **no ingress-policer**
            - **one-time-redirect**

```
                                    — one-time-redirect url rdr-url-string port
                                        port-num
                                    — no one-time-redirect
                                    — [no] shutdown
                                — vlan start [value] end [value] retail-svc-id
                                    service-id
                                — no vlan start [value] end [value]
                            — wlan-gw-group group-id
                            — no wlan-gw-group
                        — no] shutdown
```

11.25.1.3 WLAN-GW Service Commands

```
configure
    — service
        — ies service-id/vprn service-id
            — wlan-gw
                — distributed-sub-mgmt
                    — ipv6-tcp-mss-adjust segment-size
                    — no ipv6-tcp-mss-adjust
            — subscriber-interface ip-int-name
                — wlan-gw
                    — pool-manager
                        — dhcpv6-client
                            — dhcpv4-nat
                                — link-addr ipv6-address
                                — no link-addr
                                — pool-name name
                                — no pool-name
                                — [no] shutdown
                            — ia-na
                                — link-addr ipv6-address
                                — no link-addr
                                — pool-name name
                                — no pool-name
                                — [no] shutdown
                            — lease-query [max-retry Max nbr of retries]
                            — no lease-query
                            — server ipv6-address [ipv6-address... (up to 8
                                max)]
                            — no server [ipv6-address [ipv6-address... (up to 8
                                max)]]
                            — slaac
                                — link-addr ipv6-address
                                — no link-addr
                                — pool-name name
                                — no pool-name
                                — [no] shutdown
                            — source-ip ipv6-address
                            — no source-ip
                        — watermarks high high-percentage low low-percentage
```

```
                    — no watermarks
                    — wlan-gw-group nat-group-id
                    — no wlan-gw-group
                — redundancy
                    — export ip-prefix/length
                    — no export
                    — monitor ip-prefix/length
                    — no monitor
                    — [no] shutdown
            — group-interface ip-int-name [create]
            — group-interface ip-int-name [create] lns
            — group-interface ip-int-name [create] wlan
            — no group-interface ip-int-name
                — brg
                    — [no] authenticated-brg-only
                    — default-brg-profile profile-name
                    — no default-brg-profile
                    — [no] shutdown
                — dhcp
                — ip-mtu octets
                — no ip-mtu
                — sap-parameters
                    — anti-spoof {ip-mac | nh-mac}
                    — no anti-spoof
                    — description description-string
                    — no description
                    — sub-sla-mgmt
                        — def-app-profile app-profile-name
                        — no def-app-profile
                        — def-sla-profile sla-profile-name
                        — no def-sla-profile
                        — def-sub-id string sub-id
                        — def-sub-id use-auto-id
                        — no def-sub-id
                        — def-sub-profile sub-profile-name
                        — no def-sub-profile
                        — sub-ident-policy policy-name
                        — no sub-ident-policy
                — wlan-gw
                    — egress
                        — [no] agg-rate-limit
                            — rate kilobits-per-second
                            — no rate
                            — hold-time infinite
                            — hold-time [time]
                            — no hold-time
                            — qos policy-id
                            — no qos
                            — scheduler-policy scheduler-policy-name
                            — no scheduler-policy
                            — [no] shape-multi-client-only
                            — shaping {per-retailer | per-tunnel}
                            — no shaping
                    — gw-address ip-address
```

- **no gw-address**
- **gw-ipv6-address** *ipv6-address*
- **no gw-ipv6-address**
- **learn-ap-mac** [**delay-auth**]
- **no learn-ap-mac**
- **l2-access-points**
  - **l2-ap** *sap-id* [**create**]
  - **no l2-ap** *sap-id*
    - **encap-type** {**default** | **null** | **dot1q** | **qinq**}
    - **no encap-type**
    - **epipe-sap-template** name
    - **no epipe-sap-template**
    - [**no**] **shutdown**
- **l2-ap-auto-sub-id-fmt** {**include-ap-tags** | **sap-only**}
- **no l2-ap-auto-sub-id-fmt**
- **l2-ap-encap-type** {**null** | **dot1q** | **qinq**}
- **no l2-ap-encap-type**
- **mobility**
  - **hold-time** *time in s*
  - **no hold-time**
  - [**no**] **inter-vlan**
  - [**no**] **multi-tunnel-type**
  - **trigger** [**data**] [**iapp**] [**control**]
  - **no trigger**
- [**no**] **oper-down-on-group-degrade**
- **router** *router-instance*
- **no router**
- [**no**] **shutdown**
- **tcp-mss-adjust** *segment-size*
- **no tcp-mss-adjust**
- **tunnel-encaps**
  - **learn-l2tp-cookie** {**if-match** | **never** | **always**} [**cookie** *hex string*]
  - **no learn-l2tp-cookie**
- **vlan-tag-ranges**
  - **range start** [*range*] **end** [*range*]
  - **range default**
  - **no range start** [*range*] **end** [*range*]
    - **authenticate-on-dhcp**
    - **authentication**
      - **authentication-policy** *policy-name*
      - **no authentication-policy**
      - **hold-time** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
      - **vlan-mismatch-timeout** *seconds*
      - **no vlan-mismatch-timeout**
    - [**no**] **data-triggered-ue-creation**
    - **dhcp**
      - **active-lease-time** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
      - **no active-lease-time**
      - **initial-lease-time** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]

- **no initial-lease-time**
- **l2-aware-ip-address** *ip-address*
- **l2-aware-ip-address from-pool**
- **no l2-aware-ip-address**
- **primary-dns** *ip-address*
- **no primary-dns**
- **primary-nbns** *ip-address*
- **no primary-nbns**
- **secondary-dns** *ip-address*
- **no secondary-dns**
- **secondary-nbns** *ip-address*
- **no secondary-nbns**
- [**no**] **shutdown**

dhcp6

- **active-preferred-lifetime** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
- **no active-preferred-lifetime**
- **active-valid-lifetime** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
- **no active-valid-lifetime**
- **initial-preferred-lifetime** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
- **no initial-preferred-lifetime**
- **initial-valid-lifetime** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
- **no initial-valid-lifetime**
- [**no**] **shutdown**

distributed-sub-mgmt

- **accounting-policy** *policy-name*
- **no accounting-policy**
- **accounting-update-interval** [*interval*]
- **no accounting-update-interval**
- [**no**] **collect-aa-acct-stats**
- **def-app-profile**
- **def-app-profile** *profile-name*
- **no def-app-profile**
- **dsm-ip-filter** *dsm-ip-filter-name*
- **no dsm-ip-filter**
- **egress-policer** [*policer*]
- **no egress-policer**
- **ingress-policer** *policer-name*
- **no ingress-policer**
- **one-time-redirect url** *rdr-url-string* **port** *port-num*
- **no one-time-redirect**
- [**no**] **shutdown**

idle-timeout action *idle-timeout-action*
no idle-timeout
http-redirect-policy *policy-name*
no http-redirect-policy
l2-service *service-id*

```
                                        — no l2-service
                                            — description description-string
                                            — no description
                                            — [no] shutdown
                                        — nat-policy policy-name
                                        — no nat-policy
                                        — retail-svc-id service-id
                                        — [no] track-mobility

configure
    — service
        — vprn service-id
            — wlan-gw
                — distributed-sub-mgmt
                    — ipv6-tcp-mss-adjust segment-size
                    — no ipv6-tcp-mss-adjust
                — wlan-gw
                    — mobility-triggered-acct
                        — interim-update include-counters [hold-down seconds]
                        — interim-update
                        — no interim-update

configure
    — service
        — vpls service-id
            — wlan-gw
                — description description-string
                — no description
                — sap-template sap template
                — no sap-template
                — [no] shutdown
```

11.25.1.4 Data Plane Related Commands

```
config
    — isa
        — wlan-gw-group group-id [create] [redundancy unit]
        — no wlan-gw-group group-id
            — active-iom-limit number
            — no active-iom-limit
            — active-mda-limit number
            — no active-mda-limit
            — description description-string
            — no description
            — [no] distributed-sub-mgmt
                — isa-aa-group aa-group-id
                — no isa-aa-group
                — isa-aa-oversubscription-factor factor
                — no isa-aa-oversubscription-factor
            — iom slot-number type {[load-balancer] [ue-anchor]}
```

```
— no iom slot-number
— [no] mda mda-id
— nat
    — radius-accounting-policy nat-accounting-policy
    — no radius-accounting-policy
    — session-limits
        — reserved num-sessions
        — no reserved
        — watermarks high percentage low percentage
        — no watermarks
    — [no] suppress-lsn-events
    — [no] suppress-lsn-sub-blks-free
— watermarks
    — mark entity high percentage-high low percentage-low
```

11.25.1.5 RADIUS Server and Proxy Commands

```
configure
    — aaa
        — acct-on-off-group group-name [create]
        — no acct-on-off-group group-name
            — description description-string
            — no description
        — radius-server-policy policy-name [create]
        — no radius-server-policy policy-name
            — accept-script-policy policy-name
            — no accept-script-policy
            — acct-on-off monitor-group group-name
            — acct-on-off oper-state-change [group group-name]
            — no acct-on-off
            — acct-request-script-policy policy-name
            — no acct-request-script-policy
            — auth-request-script-policy policy-name
            — no auth-request-script-policy
            — [no] buffering
                — acct-interim min min-val max max-val lifetime lifetime
                — no acct-interim
                — acct-stop min min-val max max-val lifetime lifetime
                — no acct-stop
            — description description-string
            — no description
            — servers
                — access-algorithm {direct | round-robin | hash-based}
                — no access-algorithm
                — [no] disable-stickiness
                — health-check
                    — test-account
                        — interval seconds
                        — no interval
                        — password password [hash | hash2]
                        — no password
```

```
                        — [no] shutdown
                        — user-name user-name
                        — no user-name
                    — hold-down-time [sec seconds] [min minutes] [hrs hours] [days days]
                    — no hold-down-time
                    — ipv6-source-address ipv6-address
                    — no ipv6-source-address
                    — router router-instance
                    — router service-name service-name
                    — no router
                    — server server-index name server-name
                    — no server server-index
                    — source-address ip-address
                    — no source-address
                    — timeout [sec seconds] [min minutes]
                    — no timeout

configure
    — router
        — radius-server
            — server server-name [address ip-address] [secret key ] [hash | hash2] [port
                port] [create]
            — no server server-name
                — [no] accept-coa
                — acct-port port
                — no acct-port
                — auth-port port
                — no auth-port
                — coa-script-policy script-policy-name
                — no coa-script-policy
                — description description-string
                — no description
                — pending-requests-limit limit
                — no pending-requests-limit

configure
    — router
        — radius-proxy
            — server server-name [create] [purpose {[accounting | authentication]}] [wlan-
                gw-group group-id]
            — no server server-name
                — attribute-matching
                    — entry [entry] [prefix-string prefix-string] [accounting-server-
                        policy policy-name] [authentication-server-policy
                        policy-name] [suffix-string suffix-string]
                    — no entry [entry] [
                    — type [type] [vendor-id vendor-id]
                    — no type
                — cache
                    — key packet-type {accept | request} attribute-type attribute-
                        type [vendor vendor-id]
                    — no key
                    — [no] shutdown
                    — timeout  [hrs hours] [min minutes] [sec seconds]
```

```
                        — no timeout
                        — track-accounting [start] [stop] [interim-update] [accounting-
                              on] [accounting-off]
                        — no track-accounting
                        — track-authentication [accept]
                        — no track-authentication
                        — track-delete-hold-time seconds
                        — no track-delete-hold-time
                    — default-accounting-server-policy policy-name
                    — no default-accounting-server-policy
                    — default-authentication-server-policy policy-name
                    — no default-authentication-server-policy
                    — description description-string
                    — no description
                    — [no] interface interface-name
                    — load-balance-key [vendor vendor-id [vendor-id]]  attribute-type
                        attribute-type [attribute-type]
                    — load-balance-key source-ip-udp
                    — no load-balance-key
                    — python-policy name
                    — no python-policy
                    — secret secret [hash | hash2]
                    — no secret
                    — [no] send-accounting-response
                    — [no] shutdown

configure
    — service
        — vprn
            — radius-proxy
                — server server-name [create] [purpose {[accounting |
                    authentication]}] [wlan-gw-group group-id]
                — no server server-name
                    — [no] accept-coa
                    — acct-port port
                    — no  acct-port
                    — attribute-matching
                        — entry [1..32] [prefix-string prefix-string] [accounting-
                            server-policy policy-name] [authentication-server-
                            policy policy-name] [suffix-string suffix-string]
                        — no entry [1..32]
                        — type [type] [vendor-id vendor-id]
                        — no type
                    — auth-port port
                    — no auth-port
                    — cache
                        — key packet-type {accept | request} attribute-type
                            attribute-type [vendor vendor-id]
                        — no key
                        — [no] shutdown
                        — timeout [hrs hours] [min minutes] [sec seconds]
                        — no timeout
```

```
            — track-accounting [stop] [interim-update] [accounting-
                on] [accounting-off]
            — no track-accounting
        — coa-script-policy script-policy-name
        — no coa-script-policy
        — default-accounting-server-policy policy-name
        — no default-accounting-server-policy
        — default-authentication-server-policy policy-name
        — no default-authentication-server-policy
        — description description-string
        — no description
        — [no] interface interface-name
        — load-balance-key [vendor vendor-id [vendor-id]] attribute-type
                attribute-type [attribute-type]
        — load-balance-key source-ip-udp
        — no load-balance-key
        — pending-requests-limit limit
        — no pending-requests-limit
        — secret secret [hash | hash2]
        — no secret
        — [no] send-accounting-response
        — [no]shutdown
```

11.25.1.6 LUDB Matching for RADIUS Proxy Cache

```
config
    — subscriber-mgmt
        — local-user-db local-user-db-name [create]
        — no local-user-db local-user-db-name
            — dhcp
                — host
                    — match-radius-proxy-cache
                        — fail-action {continue | drop}
                        — no fail-action
                        — mac-format mac-format
                        — no mac-format
                        — match {circuit-id | mac | remote-id}
                        — match option [option] [option6 [option6]]
                        — match option6 [option]
                        — no match
                        — server [service service-id] name server-name
                        — no server
```

11.25.1.7 Port Policy Commands

```
config
    — port-policy policy-name [create]
    — no port-policy policy-name
```

— **description** *description-string*
— **no description**
— **egress-scheduler-policy** *port-sched-plcy*
— **no egress-scheduler-policy**

11.25.1.8 WIFI Aggregation and Offload – Migrant User Support Commands

configure
— **subscriber-mgmt**
— **http-redirect-policy** *policy-name* [**create**]
— **no http-redirect-policy** *policy-name*
— **description** *description-string*
— **no description**
— **dst-port** *tcp-port*
— **no dst-port**
— **forward-entries**
— **dst-ip** *ip-address* **protocol** *ip-protocol* **dst-port** *port-number*
— **dst-ip** *ip-address* **protocol** *ip-protocol* **dst-port** *port-number* **prefix-length** *prefix-length*
— **no dst-ip** *ip-address* **protocol** *ip-protocol* **dst-port** *port-number*
— **portal-hold-time** *seconds*
— **no portal-hold-time**
— **url** *rdr-url-string*
— **no url**

11.25.1.9 Show Commands

show
— **call-trace**
— **wlan-gw**
— **ue** [*ieee-address*] [**detail**]
— **router**
— **radius-proxy-server** *server-name*
— **radius-proxy-server** *server-name* **cache**
— **radius-proxy-server** *server-name* **cache hex-key** *hex-string*
— **radius-proxy-server** *server-name* **cache string-key** *string*
— **radius-proxy-server** *server-name* **cache summary**
— **radius-proxy-server** *server-name* **statistics**
— **radius-proxy-server**
— **wlan-gw**
— **isa-subnets** [**detail**]
— **isa-subnets** [**detail**] **interface** *interface-name*
— **isa-subnets prefix** *ipv6-address/prefix-length*
— **tunnel-qos** [**detail**]
— **tunnel-qos remote-ip** *ip-address* [**local-ip** *ip-address*] [**detail**]
— **tunnels local-ip** *ip-address* **remote-ip** *ip-address* ue

— **tunnels** [**local-ip** *ip-address*] [**remote-ip** *ip-address*] [**isa-group** *wlan-gw-group-id*] [**member** [1..255]] [**summary**] [**detail**]
— **tunnels** **local-ip** *ip-address* **remote-ip** *ip-address* **ue**
— **tunnels** [**local-ip** *ip-address*] [**remote-ip** *ip-address*] [**isa-group** *wlan-gw-group-id*] [**member** [1..255]] [**summary**] [**detail**]
— **tunnels** **local-ip** *ip-address* **remote-ip** *ip-address* **ue**

show
— **aaa**
— **acct-on-off-group** *group-name*
— **radius-server-policy** *policy-name* [**acct-on-off**]
— **radius-server-policy** *policy-name* **associations**
— **radius-server-policy** *policy-name* **msg-buffer-stats**
— **radius-server-policy** *policy-name* **statistics**
— **radius-server-policy** [**acct-on-off**]

show
— **isa**
— **wlan-gw-group** *wlan-gw-group-id*
— **wlan-gw-group** *wlan-gw-group-id* **associations**
— **wlan-gw-group** *wlan-gw-group-id* **member** *member-id*
— **wlan-gw-group** *wlan-gw-group-id* **member** *member-id* **resource-statistics**
— **wlan-gw-group** *wlan-gw-group-id* **member** *member-id* **statistics** [**type** *type*] [**non-zero-value-only**]
— **wlan-gw-group**

show
— **subscriber-mgmt**
— **isa-filter**
— **isa-filter** *name*
— **isa-filter** *name* **associations**
— **isa-filter** *name* **ipv4**
— **isa-filter** *name* **ipv6**
— **isa-policer** *policer-name*
— **isa-policer** *policer-name* **associations**
— **isa-policer**
— **wlan-gw**
— **ssid**
— **statistics**
— **tunnels** [**router** *router-name*] [**remote-ip** *ip-address*] [**local-ip** *ip-address*] [**encapsulation** *encap* [*encap*...(**upto 3 max**)]] [**qtag1** *qtag*] [**qtag2** *qtag*] [**ap-sap** *sap-id*] [**min-num-ue** *minimum*] [**max-num-ue** *maximum*] [**ap-mac-learn-failed** {**true** | **false**}] [**get-num-results**] [**addr-family** *family*] [**ue-type** *ue-type* [*ue-type*...(**up to 5 max**)]]
— **ue** [**vlan** *qtag*] [**mpls-label** *label*] [**retail-svc-id** *service-id*] [**ssid** *service-set-id*] [**previous-access-point** *ip-address*] [**bd** *bridge-id*]
— **ue** **mac** *ieee-address* [**bd** *bridge-id*]

11.25.1.10 Debug Commands

debug

```
— call-trace
    — [no] wlan-gw
    — group
        — statistic type type name name
        — ue ieee-address [profile trace-profile-name]
        — no ue ieee-address
```

11.25.1.11 Tools Commands

```
tools
— perform
    — aaa
        — acct-on [radius-server-policy policy-name] [force]
        — acct-off acct-off [radius-server-policy policy-name] [force] [acct-terminate-
          cause number]
— dump
    — aaa
        — radius-acct-terminate-cause
        — radius-server-policy policy-name msg-buffer [session-id acct-session-id]]
    — wlan-gw
        — isa
            — performance mda mda-id last time-span time-unit
        — ue [wlan-gw-group wlan-gw-group-id] [mda mda-id] [next-index index]
          [summary] [detail] [bd bridge-id] [ue-mac ieee-address] [ue-vlan vlan]
          [state-description state] [tunnel-router router-instance] [tunnel-source-
          ip ip-address] [tunnel-destination-ip ip-address] [tunnel-type tunnel-type]
          [ue-ip ipv4-address] [dhcp6-addr ipv6-address] [slaac-prefix ipv6-
          address] [aggregate-summary]
```

11.25.1.12 Clear Commands

```
clear
— aaa
    — radius-server-policy policy-name msg-buffer [acct-session-id acct-session-id]
    — radius-server-policy policy-name statistics [msg-buffer-only]
    — radius-server-policy policy-name server server-index statistics

clear
— router
    — wlan-gw
        — isa-subnets all
        — isa-subnets interface intfip-int-name
        — isa-subnets prefix ipv6-address/prefix-length

clear
— wlan-gw
    — isa
        — wlan-gw-group group-id member member-id resource-peak-values
```

— **wlan-gw-group** *group-id* **member** *member-id* **statistics**

11.25.2 WIFI Aggregation and Offload Commands

11.25.2.1 WIFI Aggregation and Offload Commands

11.25.2.1.1 Generic Commands

description

Syntax **description** *description-string*
no description

Context config>aaa>acct-on-off-grp
config>aaa>radius-srv-plcy
config>isa>wlan-gw-group
config>router>radius-server>server
config>router>radius-proxy>server
config>service>vprn>radius-proxy>server
config>service>vprn>radius-server>server
config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw
config>service>ies>sub-if>grp-if>sap-parameters
config>service>vprn>sub-if>grp-if>sap-parameters
config>subscr-mgmt>wlan-gw>dsm>dsm-ip-filter
config>subscr-mgmt>wlan-gw>dsm>dsm-ip-filter>ipv6
config>call-trace>trace-profile
config>subscr-mgmt>isa-filter
config>subscr-mgmt>isa-filter>ipv6>entry
config>subscr-mgmt>isa-filter>entry
config>subscr-mgmt>isa-policer

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes any description string from the context.

Parameters *description-string* — Specifies a text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

shutdown

Syntax [**no**] **shutdown**

Context config>router>radius-proxy>cache
config>router>radius-proxy>server>cache
config>router>radius-proxy>server
config>service>vprn>radius-proxy>server>cache
config>service>vprn>radius-proxy>server
config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw
config>service>ies>sub-if>wlan-gw>pool-mgr>dhcp6-client>slaac
config>service>vprn>sub-if>wlan-gw>pool-mgr>dhcp6-client>slaac
config>service>ies >sub-if>wlan-gw>pool-mgr>dhcp6-client>ia-na
config>service>vprn>sub-if>wlan-gw>pool-mgr>dhcp6-client>ia-na
config>service>ies>sub-if>wlan-gw>pool-mgr>dhcp6-client>dhcpv4-nat
config>service>vprn>sub-if>wlan-gw>pool-mgr>dhcp6-client>dhcpv4-nat
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>brg
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>brg
config>service>ies>sub-if>grp-if>brg
config>service>vprn>sub-if>grp-if>brg

Description This command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command.

The **shutdown** command administratively disables an entity. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they can be deleted.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system generated configuration files.

The **no** form of this command puts an entity into the administratively enabled state.

Default no shutdown

subscriber-mgmt

Syntax **subscriber-mgmt**

Context config

Description This command enables the context to configure subscriber management entities. A subscriber is uniquely identified by a subscriber identification string. Each subscriber can have several DHCP sessions active at any time. Each session is referred to as a subscriber host and is identified by its IP address and MAC address.

All subscriber hosts belonging to the same subscriber are subject to the same hierarchical QoS (HQoS) processing. The HQoS processing is defined in the sub-profile (the subscriber profile). A sub-profile refers to an existing scheduler policy (configured in **the config>qos>scheduler-policy** context) and offers the possibility to overrule the rate of individual schedulers within this policy.

Because all subscriber hosts use the same scheduler policy instance, they must all reside on the same complex.

11.25.2.1.2 WLAN-GW Commands

Note: The **wlan-gw** commands apply only to the 7750 SR platform.

wlan-gw

Syntax [**no**] **wlan-gw**

Context config>subscr-mgmt
config>router
config>service>vprn

Description This command enables the context to configure WLAN Gateway parameters.

distributed-sub-mgmt

Syntax **distributed-sub-mgmt**

Context config>subscr-mgmt>wlan-gw
config>router>wlan-gw
config>service>vprn>wlan-gw

Description This command enables the context to configure profiles, templates and policies that can be applied to DSM subscribers.

virtual-chassis-identifier

Syntax **virtual-chassis-identifier** *dual-homing-key*
no virtual-chassis-identifier

Context config>subscr-mgmt>wlan-gw

Description This command specifies a virtual chassis identifier that can link two wlan-gw's together.

The **no** form of this command removes the dual-homing-key

Parameters *dual-homing-key* — Specifies the name of the dual homing key up to 16 characters.

ipv6-tcp-mss-adjust

Syntax **ipv6-tcp-mss-adjust** *segment-size*
no ipv6-tcp-mss-adjust

Context config>router>wlan-gw>dsm
config>service>vprn>wlan-gw>dsm

Description This command specifies the value used for TCP-MSS-adjust in the IPv6 upstream direction for DSM. The downstream direction for both IPv4 and IPv6 are both configured under the group-interface. The upstream direction for IPv4 NAT hosts is configured under the NAT policy.

The defined segment size is inserted in a TCP SYN message if there is no existing MSS option or the value in the MSS option is bigger than the configured value.

The **no** form of this command disables upstream TCP MSS adjust for IPv6 DSM.

Parameters *segment-size* — Specifies the segment size to be inserted.

Values 160 to 10240

gtp-peer-clear-timeout

Syntax **gtp-peer-clear-timeout** *seconds*
no gtp-peer-clear-timeout

Context config>service>vprn>wlan-gw>dsm

Description This command configures a GTP peer cleanup timeout to terminate a handover wait state.

Parameters *seconds* — Specifies a GTP peer cleanup timeout, in seconds, to terminate a handover wait state.

Values 0 to 3600

mobility-triggered-acct

Syntax **mobility-triggered-acct**

Context config>router>wlan-gw
config>service>vprn>wlan-gw

Description This command enters the configuration context of mobility-triggered-accounting in wlan-gw context under router or VPRN service.

interim-update

Syntax **interim-update**
interim-update include-counters [**hold-down** *seconds*]
no interim-update

Context config>router>wlan-gw>mobility-triggered-acct
config>service>vprn>wlan-gw>mobility-triggered-acct

Description This command enables the inclusion of counters with a **hold-down** time option in mobility-triggered interim-updates. When enabled, to disable the inclusion of counters, interim updates must be disabled and then re-enabled without the **include-counters** keyword. By default, the **hold-down** time is not imposed.

The **no** form of this command disables generation of flash interim accounting updates to RADIUS when change in location of the UE is detected.

Parameters **include-counters** — Specifies the inclusion of counters in mobility triggered interim-updates.

seconds — Specifies the time, in seconds, that must elapse after a mobility- triggered interim with counters sent for the next mobility-triggered interim with counters to be sent.

Values 60 to 864000

11.25.2.1.3 ISA Commands

isa-filter

Syntax **isa-filter** *name* [**type** {**dsm**}] [**create**]
no isa-filter *name*

Context config>subscr-mgmt

Description This command enables the context to configure ISA filter parameters.

Parameters *name* — Specifies the name of the filter

type dsm — Selects DSM as the type.

default-action

Syntax **default-action** {**drop** | **forward**}
no default-action

Context config>subscr-mgmt>isa-filter
config>subscr-mgmt>isa-filter>ipv6

Description This command specifies what should happen to packets that do not match any of the configured entries.

The **no** form of this command reverts to the default value.

Default default-action drop

Parameters **drop** — Specifies that packets matching the filter entry are dropped.

forward — Specifies that packets matching the filter entry are forwarded.

entry

Syntax **entry** *entry-id* [**create**]
no entry *entry-id*

Context config>subscr-mgmt>isa-filter
config>subscr-mgmt>isa-filter>ipv6

Description This command creates a new entry for this filter. When processing a packet, entries are matched in order, starting with the lowest entry-id. A maximum of 128 IPv4 and 128 IPv6 DSM filter entries are allowed.

The **no** form of this command removes the specified entry from the ISA filter.

Parameters *entry-id* — Specifies the numeric identifier for the filter entry.

action

Syntax **action** {**drop** | **forward** | **none**}
no action

Context config>subscr-mgmt>isa-filter>entry
config>subscr-mgmt>isa-filter>ipv6>entry

Description This command specifies what should happen to packets that do match this entry. If the configured action is none, this entry is not applied and processing continues to match against subsequent entries.

The **no** form of this command reverts to the default value.

Default action none

Parameters **drop** — Drops the packet.

forward — Forwards the packet.

none — Disables this entry. The packet processing continues with the next entry.

match

Syntax **match protocol** {**any** | **icmp** | **tcp** | **upd** | **gre**}
no match

Context config>subscr-mgmt>entry
config>subscr-mgmt>ipv6>entry

Description This command creates a match context for this entry. The *protocol* value specifies which Layer-4 protocol the packet should match.

The **no** form of this command removes the match context of this entry.

Default match protocol any

Parameters *protocol* — Specifies that the only supported match context is **protocol**.

any — Specifies to match any protocol.

icmp — Specifies to match ICMP packets in a v4 filter.

tcp — Specifies to match TCP packets.

udp — Specifies to match UDP packets.

gre — Specifies to match GRE over IP packets.

dst-port

Syntax **dst-port** *operator port-number*
no dst-port

Context config>subscr-mgmt>isa-filter>entry>match
config>subscr-mgmt>isa-filter>ipv6>entry>match

Description This command specifies that the packet's UDP/TCP dst-port must match a specific value. This command is not valid in a match context that is not specific to UDP or TCP.

The **no** form of this command removes matching of the layer-4 port.

Parameters *operator* — Specifies that the only supported value is eq (equal to). The destination port value must be equal to the *port-number* value.

port-number — Specifies the number of the port to match.

Values 0 to 65535

dst-ip

Syntax **dst-ip** *ip-prefix/length*
no dst-ip

Context config>subscr-mgmt>isa-filter>entry
config>subscr-mgmt>isa-filter>ipv6>entry

Description This command specifies that the packet's destination IP address must match the specified IP prefix and mask.

The **no** form of this command disables the match on the destination IP.

Parameters *ip-prefix/length* — Specifies the IP prefix to match.

src-ip

Syntax **src-ip** *ip-prefix/length*
no src-ip

Context config>subscr-mgmt>isa-filter>entry>match
config>subscr-mgmt>isa-filter>ipv6>entry>match

Description This command configures the source IP or IPv6 address match condition.

The **no** form of this command reverts to the default value.

src-port

Syntax **src-port** *operator port-number*
no src-port

Context config>subscr-mgmt>isa-filter>entry>match
config>subscr-mgmt>isa-filter>ipv6>entry>match

Description This command configures the source port match condition.

The **no** form of this command reverts to the default value.

cbs

Syntax **cbs** *burst-size*
no cbs

Context config>subscr-mgmt>isa-policer

Description This command specifies the committed burst-size value of this policer. This can only be set on dual-bucket-bandwidth policers.

The **no** form of this command reverts to its default.

Default cbs 0

Parameters *burst-size* — Specifies the committed burst-size in kbytes.

Values 0 to 131071

mbs

Syntax **mbs** *burst-size*
no mbs

Context config>subscr-mgmt>isa-policer

Description This command specifies the maximum burst-size value of this policer.

The **no** form of this command reverts to its default.

Default mbs 0

Parameters *burst-size* — The maximum burst-size in kbytes.

Values 0 to 131071

rate

Syntax **rate** *rate* [**cir** *cir-rate*]
no rate

Context config>subscr-mgmt>isa-policer

Description This command specifies at which rate the policer drains packets. The cir value is only supported on dual-bucket-bandwidth policers. If rate **max** is configured, no actual rate limitations are applied.

The **no** form of this command reverts to the default.

Parameters *rate* — Specifies the rate in Kbps.

Values 1 to 100000000, **max**

Default max

cir-rate — Specifies the CIR rate in Kbps.

Values 1 to 100000000, **max**

Default max

isa-policer

Syntax **isa-policer** *policer-name* [**type** *policer-type*] [**create**]
no isa-policer *policer-name*

Context config>subscr-mgmt

Description This command creates the context to configure an ISA policer. When creating a policer for the first time, both the **create** and **type** parameters are required.

The **no** form of this command removes the.

Parameters *policer-name* — Specifies the name by which this policer is referenced up to 32 characters.

policer-type — Specifies the policer type. The dual-bucket-bandwidth policer applies both a CIR and PIR.

Values single-bucket-bandwidth, dual-bucket-bandwidth

action

Syntax **action** {**permit-deny** | **priority-mark**}
no action

Context config>subscr-mgmt>isa-policer

Description This command specifies what happens to packets that are in-profile and out-of-profile.

The **no** form of this command reverts to the default value.

Default action permit-deny

Parameters **permit-deny** — Drops all packets that are out of profile (they do not conform to the PIR).

priority-mark — Currently not supported. The policer will take no action.

adaptation-rule

Syntax **adaptation-rule pir** *adaptation-rule* [**cir** *adaptation-rule*]
no adaptation-rule

Context config>subscr-mgmt>isa-policer

Description For operational efficiency, the operational rate of a policer cannot take on every value in the configurable range. This configuration defines a rule that must be followed when mapping a configured rate to an operational rate.

The cir **adaptation-rule** can only be set on dual-bucket-bandwidth policers.

The **no** form of this command reverts to its default.

Default adaptation-rule pir closest cir closest

Parameters **pir** *adaptation-rule* — Configures the rules to compute the PIR operational rates.

Values **min** — Specifies that the operational rate must minimally be the configured rate. The first operational value bigger or equal to the configured rate is chosen.

max — Specifies that the operational rate may maximally be the configured rate. The first operational value smaller or equal to the configured rate is chosen.

closest — Chooses the operational value closest to the configured value, lower or higher.

cir *adaptation-rule* — Configures the rules to compute the CIR operational rates.

Values adaptation-rule

11.25.2.1.4 RADIUS Server Policy Commands

acct-on-off-group

Syntax **acct-on-off-group** *group-name* [**create**]
no acct-on-off-group *group-name*

Context config>aaa

Description This command creates an acct-on-off-group.

An acct-on-off-group can be referenced by:

- A single radius-server-policy as controller — The acct-on-off oper-state of the acct-on-off-group is set to the acct-on-off oper-state of the radius-server-policy (acts as master).

- Multiple radius-server-policies as monitor — The acct-on-off oper-state of the radius-server-policy is inherited from the acct-on-off oper-state of the acct-on-off group. (acts as a slave).

The **no** form of this command deletes the acct-on-off-group.

Parameters *group-name* — Specifies the name of an acct-on-off group up to 32 characters.

radius-server-policy

Syntax **radius-server-policy** *policy-name* [**create**]
no radius-server-policy *policy-name*

Context config>aaa

Description This command creates a radius-server-policy.

A RADIUS server policy can be used in

- radius-proxy, for application like EAP authentication for WIFI access
- authentication policy, for Enhanced Subscriber Management authentication
- radius accounting policy, for Enhanced Subscriber Management accounting
- dynamic data service RADIUS accounting
- AAA route downloader

The **no** form of this command removes the policy name from the configuration.

Parameters *policy-name* — Specifies the name of the radius-server-policy up to 32 characters.

create — Keyword used to create a radius-server-policy name. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

accept-script-policy

Syntax **accept-script-policy** *policy-name*
no accept-script-policy

Context config>aaa>radius-srv-plcy

Description This command specifies name of the radius-script-policy to be applied for access-accept.

Parameters *policy-name* — Specifies the name of the accept-script-policy up to 32 characters.

acct-on-off

Syntax **acct-on-off**

acct-on-off monitor-group *group-name*
acct-on-off oper-state-change [**group** *group-name*]

Context config>aaa>radius-srv-plcy

Description This command controls the sending of Accounting-On and Accounting-Off messages and the acct-on-off oper-state of the radius-server-policy:

acct-on-off: enables the sending of Accounting-On and Accounting-Off messages for this radius-server-policy. The acct-on-off oper-state is always not blocked.

acct-on-off oper-state-change [group *group-name*]: enables the sending of Accounting-On and Accounting-Off messages for this radius-server-policy. The acct-on-off oper-state is function of the Accounting-response received for the Accounting-On and Accounting-Off. Optionally, sets the acct-on-off oper-state of the acct-on-off-group.

acct-on-off monitor-group *group-name*: no Accounting-On and Accounting-Off messages are sent for this radius-server-policy. The acct-on-off oper-state is inherited from the acct-on-off-group.

The **no** form of this command disables the sending of Accounting-On and Accounting-Off messages.

Parameters *group-name* — Specifies the name of an acct-on-off group up to 32 characters.

acct-request-script-policy

Syntax **acct-request-script-policy** *policy-name*
no acct-request-script-policy

Context config>aaa>radius-srv-plcy

Description This command specifies the name of the acct-request-script-policy pointing to the Python script to be applied for RADIUS accounting request messages.

Parameters *policy-name* — Specifies the name of the acct-request-script-policy up to 32 characters.

auth-request-script-policy

Syntax **auth-request-script-policy** *policy-name*
no auth-request-script-policy

Context config>aaa>radius-srv-plcy

Description This command specifies the name of the auth-request-script-policy pointing to the Python script to be applied for RADIUS access request messages.

Parameters *policy-name* — Specifies the name of the auth-request-script-policy up to 32 characters

buffering

Syntax [**no**] **buffering**

Context config>aaa>radius-srv-plcy

Description This command enables the context to configure RADIUS message buffering.

The **no** form of this command disables RADIUS message buffering.

acct-interim

Syntax **acct-interim min** *min-val* **max** *max-val* **lifetime** *lifetime*
no acct-interim

Context config>aaa>radius-srv-plcy>servers>buffering

Description This command enables RADIUS accounting interim update message buffering.

1. The message is stored in the buffer, a lifetime timer is started and the message is sent to the RADIUS server
2. If after retry*timeout seconds no RADIUS accounting response is received for the interim update then a new attempt to send the message is started after minimum[(min-val*2n), max-val] seconds.
3. Repeat step 2 until for one of the following:
    a. a RADIUS accounting response is received.
    b. the lifetime of the buffered message expires.
    c. a new RADIUS accounting interim-update or a RADIUS accounting stop for the same accounting session-id and radius-server-policy is stored in the buffer.
    d. the message is manually purged from the message buffer via a clear command.
4. The message is purged from the buffer.

The **no** form of this command disables RADIUS accounting interim update message buffering.

Parameters *min-val* — Specifies the minimum interval in seconds between attempts to resend the RADIUS accounting interim update.

Values 1 to 3600

max-val — Specifies the maximum interval in seconds between attempts to resend the RADIUS accounting interim update.

Values 1 to 3600

lifetime — Specifies the lifetime in hours.

Values 1 to 25

acct-stop

Syntax **acct-stop min** *min-val* **max** *max-val* **lifetime** *lifetime*
no acct-stop

Context config>aaa>radius-srv-plcy>servers>buffering

Description This command enables RADIUS accounting stop message buffering.

1. The message is stored in the buffer, a lifetime timer is started and the message is sent to the RADIUS server
2. If after retry*timeout seconds no RADIUS accounting response is received for the accounting stop, then a new attempt to send the message is started after minimum[(min-val*2n), max-val] seconds.
3. Repeat step 2 until
   a. a RADIUS accounting response is received, or
   b. the lifetime of the buffered message expires, or
   c. the message is manually purged from the message buffer via a clear command
4. The message is purged from the buffer.

The **no** form of this command disables RADIUS accounting stop message buffering.

Parameters *min-val* — Specifies the minimum interval in seconds between attempts to resend the RADIUS accounting stop.

Values 1 to 3600

max-val — Specifies the maximum interval in seconds between attempts to resend the RADIUS accounting stop.

Values 1 to 3600

lifetime — Specifies the lifetime in hours.

Values 1 – 25

servers

Syntax **servers**

Context config>aaa>radius-srv-plcy

Description This command enables the context to configure radius-server-policy parameters.

access-algorithm

Syntax **access-algorithm** {**direct** | **round-robin** | **hash-based**}
no access-algorithm

Context config>aaa>radius-srv-plcy>servers

Description This command configures the algorithm used to select a RADIUS server from the pool of configured RADIUS servers.

Default access-algorithm direct

Parameters **direct** — Specifies that the first server is used as primary server for all requests, the second as secondary and so on.

round-robin — Specifies that the first server is used as primary server for the first request, the second server as primary for the second request, and so on. If the router gets to the end of the list, it starts again with the first server.

hash-based — Select a RADIUS server based on the calculated hash result of the configured **load-balance-key** under the **radius-proxy server** hierarchy. This parameter is only applicable for radius-proxy server scenarios and results in an unpredictable RADIUS server selection if used in other scenarios.

disable-stickiness

Syntax [**no**] **disable-stickiness**

Context config>aaa>radius-srv-plcy>servers

Description This command disables a subscriber RADIUS accounting session from sticking with a single server under normal working conditions. If a direct algorithm is used, all subscriber RADIUS sessions will go directly to the server with the lowest configured index. If a failure occurs, a new in-service server with the next lowest index is used. When the original server recovers, if stickiness is not disabled, all existing sessions will continue to use the new server. This command disables stickiness, and as a result, the recovered original RADIUS server will again service every subscriber. If a round-robin algorithm is used and stickiness is not disabled, an accounting session for a particular subscriber (or host, depending on the accounting mode) will stay with the same server. This command removes the stickiness and all subscriber accounting messages will go through the list of servers in a round-robin manner.

health-check

Syntax **health-check**

Context config>aaa>radius-server-policy>servers

Description This command enables the context to configure health check parameters for the RADIUS server.

test-account

Syntax **test-account**

Context config>aaa>radius-srv-plcy>servers>health-check

Description This command sets up a test account as a probing mechanism to check the connectivity of all configured RADIUS authentication servers within the RADIUS server policy.

interval

Syntax **interval** *seconds*
no interval

Context config>aaa>radius-srv-plcy>servers>health-check>test-account

Description This command specifies the intervals at which the test account will send its access requests to probe the RADIUS servers.

Default interval 3

Parameters *seconds* — Specifies the probing interval.

Values 1 to 60

password

Syntax **password** *password* [**hash** | **hash2**]
no password

Context config>aaa>radius-srv-plcy>servers>health-check>test-account
config>subscr-mgmt>vrgw>brg>brg-profile>radius-authentication

Description This command specifies the password that the test account will use to send access requests to probe the RADIUS servers.

Parameters *password* — Specifies the probing password up to 64 characters

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

shutdown

Syntax [**no**] **shutdown**

Context config>aaa>radius-srv-plcy>servers>health-check>test-account

Description This command disables the test account that probes the RADIUS server.

The **no** form of this command enables the capability.

Default shutdown

user-name

Syntax **user-name** *user-name*
no user-name

Context config>aaa>radius-srv-plcy>servers>health-check>test-account

Description This command specifies the username that the test account will use to send its access requests to probe the RADIUS servers.

The no form of this command removes the username from the test-account configuration.

Parameters *user-name* — Specifies the probing username up to 64 characters.

retry

Syntax **retry** *count*
no retry

Context config>aaa>radius-srv-plcy>servers

Description This command configures the number of times the router attempts to contact the RADIUS server, if not successful the first time.

The **no** form of this command reverts to the default.

Default retry 3

Parameters *count* — Specifies the number of times a signaling request message is transmitted towards the same peer.

Values 1 to 256

router

Syntax **router** *router-instance*
router service-name *service-name*
no router

Context config>aaa>radius-srv-plcy>servers

Description This command specifies the virtual router instance applicable for the set of configured RADIUS servers. This value cannot be changed once a RADIUS server is configured for this policy.

The **no** form of this command reverts to the default.

Parameters *router-instance* — Specifies the router instance.

Values

service-name	Service name up to 64 characters.	
	router-instance:	router-name, service-id
	router-name:	Base, management
	service-id:	1 to 2147483647

service-name — Specifies the router name service-id up to 64 characters.

server

Syntax **server** *server-index* **name** *server-name*
no server *server-index*

Context config>aaa>radius-srv-plcy>servers

Description This command adds a RADIUS server.

The **no** form of this command removes a RADIUS server.

Parameters *index* — The index for the RADIUS server. The index determines the sequence in which the servers are queried for authentication requests. Servers are queried in order from lowest to highest index.

Values 1 to 5

server-name — Specifies the server name up to 32 characters.

source-address

Syntax **source-address** *ip-address*
no source-address

Context config>aaa>radius-srv-plcy>servers

Description This command configures the source address of the RADIUS packet. The system IP address must be configured in order for the RADIUS client to work. See Configuring a System Interface in the 7750 SR OS Configuration Guide.

Note: The system IP address must only be configured if the source-address is not specified. When the no source-address command is executed, the source address is determined at the moment the request is sent. This address is also used in the nas-ip-address attribute: over there it is set to the system IP address if no source-address was given.

The **no** form of this command reverts to the default value.

Parameters *ip-address* — Specifies the source address of radius packet.

timeout

Syntax **timeout** [**sec** *seconds*] [**min** *minutes*]
no timeout

Context config>aaa>radius-srv-plcy>servers

Description This command configures the time the router waits for a response from a RADIUS server.

The no form of this command reverts to the default value.

Default timeout sec 5

Parameters *seconds* — Specifies the number of seconds for the timeout.

Values 1 to 59

minutes — Specifies the number of minutes for the timeout.

Values 1 to 5

Values Max. value = 5 min 40 sec

hold-down-time

Syntax **hold-down-time** [**sec** *seconds*] [**min** *minutes*] [**hrs** *hours*] [**days** *days*]
no hold-down-time

Context config>aaa>radius-srv-plcy>servers

Description This command configures the hold time before re-using a RADIUS server.

The **no** form of this command reverts to the default value.

Default hold-down-time sec 30

Parameters *seconds* — Specifies the number of seconds for the hold down time.

Values 1 to 59

minutes — Specifies the number of minutes for the hold down time.

Values 1 to 59

hours — Specifies the number of hours for the hold down time.

Values 1 to 23

days — Specifies the number of days for the hold down time.

Values 1 to 1

ipv6-source-address

Syntax **ipv6-source-address** *ipv6-address*
no ipv6-source-address

Context config>aaa>radius-srv-plcy>servers

Description This command configures the source address of an IPv6 RADIUS packet.

When no ipv6-source-address is configured, the system IPv6 address (inband RADIUS server connection) or Boot Option File (BOF) IPv6 address (outband RADIUS server connection) must be configured in order for the RADIUS client to work with an IPv6 RADIUS server.

This address is also used in the NAS-IPv6-Address attribute.

The **no** form of this command reverts to the default value.

Parameters *ipv6-address* — Specifies the source address of an IPv6 RADIUS packet.

11.25.2.2 CLI Command Description for RADIUS Server

radius-server

Syntax **radius-server**

Context config>router
config>service>vprn

Description This command enters the radius-server configuration context under router or VPRN service.

server

Syntax **server** *server-name* [**address** *ip-address*] [**secret** *key*] [**hash** | **hash2**] [**create**]
no server *server-name*

Context config>router>radius-server
config>service>vprn>radius-server

Description This command either specifies an external RADIUS server in the corresponding routing instance or enters configuration context of an existing server. The configured server could be referenced in the radius-server-policy.

The **no** form of this command removes the parameters from the server configuration.

Parameters *server-name* — Specifies the name of the external RADIUS server.

ip-address — Specifies the IPv4 or IPv6 IP address of the external RADIUS server.

key — Specifies the shared secret key of the external RADIUS server, up to 64 characters.

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

accept-coa

Syntax [**no**] **accept-coa**

Context config>router>radius-server>server
config>service>vprn>radius-server>server

Description This command configures this server for Change of Authorization messages. The system will process the CoA request from the external server if configured with this command; otherwise the CoA request is dropped.

The **no** form of this command disables the command.

acct-port

Syntax **acct-port** *port*
no acct-port

Context config>router>radius-server>server
config>service>vprn>radius-server>server

Description This command specifies the UDP listening port for RADIUS accounting requests.

The **no** form of this commands resets the UDP port to its default value (1813)

Default acct-port 1813

Parameters *port* — Specifies the UDP listening port for accounting requests of the external RADIUS server.

Values 1 to 65535

auth-port

Syntax **auth-port** *port*
no auth-port

Context config>router>radius-server>server
config>service>vprn>radius-server>server

Description This command specifies the UDP listening port for RADIUS authentication requests.

The **no** form of this commands resets the UDP port to its default value (1812)

Default auth-port 1812

Parameters *port* — Specifies the UDP listening port for accounting requests of the external RADIUS server.

Values 1 to 65535

coa-script-policy

Syntax **coa-script-policy** *policy-name*
no coa-script-policy

Context config>router>radius-server>server
config>service>vprn>radius-server>server

Description This command specifies radius-script-policy for CoA-Request sent from this RADIUS server.

The **no** form of this command removes the policy name from the configuration.

Parameters *policy-name* — Specifies the name of radius-script-policy up to 80 characters.

pending-requests-limit

Syntax **pending-request-limit** *limit*
no pending-request-limit

Context config>router>radius-server>server
config>service>vprn>radius-server>server

Description This command specifies the per-server maximum number of outstanding requests sent to the RADIUS server. If the maximum number is exceeded, the next RADIUS server in the pool is selected.

The **no** form of this command removes the limit value from the configuration.

Default pending-requests-limit 4096

Parameters *limit* — Specifies the maximum number of outstanding requests sent to the RADIUS server.

Values 1 to 4096

11.25.2.3 CLI Command Description for RADIUS Proxy Server

radius-proxy

Syntax **radius-proxy**

Context config>router
config>service>vprn

Description This command context to configure RADIUS proxy parameters.

server

Syntax **server** *server-name* [**create**] [**purpose** {[**accounting** | **authentication**]}] [**wlan-gw-group** *group-id*]
no server *server-name*

Context config>router>radius-proxy
config>service>vprn>radius-proxy

Description This command creates a RADIUS-proxy server in the corresponding routing instance. The proxy server can be configured for the purpose of proxying authentication or accounting or both.

If a WLAN-GW ISA group is specified, then the RADIUS proxy server is instantiated on the set of ISAs in the specified wlan-gw group. The RADIUS messages from the AP are load-balanced to these ISAs. The ISA that processes the RADIUS message then hashes this message to the ISA that anchors the UE. The hash is based on UE MAC address (required to be present in the calling-station-id attribute) in the RADIUS message.

If the **create** parameter is not specified, then this command enters configuration context of the specified RADIUS-proxy server.

The **no** form of this command removes the server-name and parameters from the radius-proxy configuration.

Default purpose authentication

Parameters *server-name* — Specifies the name of the RADIUS-proxy server.

create — Specifies that the system will create the specified RADIUS-proxy server.

purpose — Specifies the purpose the RADIUS-proxy server.

Values accounting — proxy accounting packets
authentication — proxy authentication packets
both — Specifies both accounting and authentication proxy accounting packets

group-id — Specifies the WLAN-GW ISA group.

attribute-matching

Syntax **attribute-matching**

Context config>router>radius-proxy>server
config>service>vprn>radius-proxy>server

Description This command enables the context for selecting the RADIUS policy for authentication and accounting based on the RADIUS attribute. This feature is supported for both the ESM RADIUS proxy and the ISA RADIUS proxy.

entry

Syntax **entry** [*entry*] [**prefix-string** *prefix-string*] [**accounting-server-policy** *policy-name*] [**authentication-server-policy** *policy-name*] [**suffix-string** *suffix-string*]
no entry [*entry*]

Context config>router>radius-proxy>server>attribute-matching
config>service>vprn>radius-proxy>server>attribute-matching

Description This command matches the specified prefix or suffix string with the selected accounting server policy or authentication server policy.

Parameters *entry* — Specifies an entry ID.

Values 1 to 32

prefix-string — Specifies the prefix string for matching up to 128 characters. If the suffix-string is also used, the combined length cannot exceed 126 characters.

suffix-string — Specifies the suffix string for matching up to 126 characters. If the prefix-string is also used, the combined length cannot exceed 126 characters.

policy-name — Specifies the RADIUS accounting or authentication policy up to 32 characters.

type

Syntax **type** [*type*] [**vendor-id** *vendor-id*]
no type

Context config>router>radius-proxy>server>attribute-matching
config>service>vprn>radius-proxy>server>attribute-matching

Description This command specifies the RADIUS VSA type for the entries to be matched with.

Parameters type — Specifies the RADIUS server policy matching attribute-type and vendor-id.

Values 1 to 255

vendor-id — Specifies the vendor ID number.

Values 1 to 16777215, **nokia**

cache

Syntax **cache**

Context config>router>radius-proxy>server
config>service>vprn>radius-proxy>server

Description This command enters the cache configuration context under radius-proxy server. The cache contains per-subscriber authentication information learned from RADIUS authentication messages, and is used to authorize subsequent DHCP requests.

default-accounting-server-policy

Syntax **default-accounting-server-policy** *policy-name*
no default-accounting-server-policy

Context config>router>radius-proxy>server

config>service>vprn>radius-proxy>server

Description This command specifies the default radius-server-policy for RADIUS accounting. This policy is used when there is no specific match based on username.

The **no** form of this command removes the policy name from the configuration.

Parameters *policy-name* — Specifies the name of the default RADIUS server policy associated with this RADIUS Proxy server for accounting purposes.

default-authentication-server-policy

Syntax **default-authentication-server-policy** *policy-name*
no default-authentication-server-policy

Context config>router>radius-proxy>server
config>service>vprn>radius-proxy>server

Description This command specifies the default radius-server-policy for RADIUS authentication. This policy is used when there is no specific match based on username.

The **no** form of this command removes the policy name from the configuration.

Parameters *policy-name* — Specifies the name of the default RADIUS server policy associated with this RADIUS proxy server for authentication purposes.

interface

Syntax [**no**] **interface** *ip-int-name*

Context config>router>radius-proxy>server
config>service>vprn>radius-proxy>server

Description This command configures the IP interface the RADIUS-proxy server will bind to. One RADIUS-proxy server could bind to multiple interfaces.

Parameters *ip-int-name* — Specifies the name of an IP interface.

load-balance-key

Syntax **load-balance-key** [**vendor** *vendor-id* [*vendor-id*]] **attribute-type** *attribute-type* [*attribute-type*]
load-balance-key source-ip-udp
no load-balance-key

Context config>router>radius-proxy>server
config>service>vprn>radius-proxy>server

Description This command specifies the key used in calculating a hash to select an external RADIUS server from the pool of configured servers.

The key can be the source IP and source UDP port tuple, or the specified RADIUS attribute in RADIUS packets.

The **no** form of this command removes the parameters from the configuration.

Parameters *vendor-id* — Specifies the vendor-id of vendor-specific attribute

Values 0 to 16777215

attribute-type — Specifies that the key is constructed with the attributes in the RADIUS message.

Values 1 to 255

source-ip-udp — Specifies that the key consists of the source IP address and source UDP port of the RADIUS message.

python-policy

Syntax **python-policy** *name*
no python-policy

Context config>router>radius-proxy>server

Description This command specifies the Python policy used to change the RADIUS attributes of the different RADIUS messages.

The **no** form of this command removes the name from the configuration.

Parameters *name* — Specifies the Python policy name up to 32 characters.

secret

Syntax **secret** *secret* [**hash** | **hash2**]
no secret

Context config>router>radius-proxy>server
config>service>vprn>radius-proxy>server

Description This command configures the shared secret key. The RADIUS client must have the same key to communicate with the RADIUS-proxy server.

The **no** form of this command removes the parameters from the configuration.

Parameters **secret** *key* — Specifies the secret key up to 64 characters to access the RADIUS server. This secret key must match the password on the RADIUS server.

Values hash-key: Up to 33 characters

hash2-key: Up to 55 characters.

hash — Specifies that the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified,

hash2 — Specifies that the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

send-accounting-response

Syntax [**no**] **send-accounting-response**

Context config>router>radius-proxy>server
config>service>vprn>radius-proxy>server

Description This command results in the system to always generate RADIUS accounting-response to acknowledge RADIUS accounting-request received from the RADIUS client.

The **no** form of this command disables the command.

key

Syntax **key packet-type** {**accept** | **request**} **attribute-type** *attribute-type* [**vendor** *vendor-id*]
no key

Context config>router>radius-proxy>server>cache
config>service>vprn>radius-proxy>server>cache

Description This command specifies the RADIUS cache key that is used to match the information in subsequent DHCP requests for authorization.

Parameters **packet-type** — Specifies the packet type of the RADIUS messages to use to generate the key for the cache of this RADIUS proxy server.

Values accept, request

attribute-type — Specifies the RADIUS attribute type to cache for this RADIUS proxy. server.

Values 1 to 255

vendor *vendor-id* — Specifies the RADIUS vendor ID.

Values 1 to 16777215, nokia

timeout

Syntax **timeout** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no timeout

Context config>router>radius-proxy>server>cache
config>service>vprn>radius-proxy>server>cache

Description This command configures the time for which the cache entry is kept if there is no corresponding DHCP DISCOVER. At the expiry of this time, the cache entry is deleted.

The **no** form of this command reverts to the default value.

Default timeout min 5

Parameters *hours* — Specifies, in hours, the timeout after which an entry in the cache will expire.

Values 1

minutes — Specifies, in minutes, the timeout after which an entry in the cache will expire.

Values 1 to 59

seconds — Specifies, in seconds, the timeout after which an entry in the cache will expire.

Values 1 to 59

track-accounting

Syntax **track-accounting** [**start**] [**stop**][**interim-update**][**accounting-on**] [**accounting-off**]
no track-accounting

Context config>router>radius-proxy>server>cache
config>service>vprn>radius-proxy>server>cache

Description This command specifies the type of RADIUS accounting packets from RADIUS client (a WIFI AP) that the router should track.

The **no** form of this command removes the parameters from the configuration.

Parameters **start** — Specifies that the router will update the associated ESM-host with the RADIUS client (for example, a WIFI AP) that generated the accounting-start. This is required in cases where a UE roams to a new AP that does not re-authenticate due to key caching.

stop — Specifies that the router will remove the corresponding ESM host and forward the accounting-stop packet to the external RADIUS server.

accounting-on | accounting-off — Specifies that the router will remove all ESM hosts associated with the RADIUS client (a WIFI AP), and forward the accounting-on packet to the external RADIUS server.

interim-update — Specifies that the router will update the associated ESM-host with the RADIUS client (a WIFI AP) that generated the interim-update. The interim-updates with the updated information are sent to the RADIUS server as scheduled.

track-authentication

Syntax **track-authentication** [**accept**]
no track-authentication

Context config>router>radius-proxy>server>cache
config>service>vprn>radius-proxy>server>cache

Description This command specifies if RADIUS authentication (from the AP) should be tracked in order to update the ESM host with the RADIUS client (for example, WIFI AP) on UE mobility. It also specifies the authentication packet from RADIUS client (for example, a WIFI AP) that the router should track for mobility.

The **no** form of this command stops tracking authentication for UE mobility.

Default track-authentication accept

Parameters **accept** — Indicates access-accept is tracked for mobility.

track-delete-hold-time

Syntax **track-delete-hold-time** *seconds*
no track-delete-hold-time

Context config>router>radius-proxy>server>cache

Description This command specifies the delete hold-time in case the DHCP host gets a trigger to delete from the matched RADIUS Proxy server.

The **no** form of this command reverts to the default.

Default track-delete-hold-time 0

Parameters *seconds* — Specifies the delete hold time, in seconds.

Values 0 to 600

11.25.2.4 LUDB Matching of RADIUS Proxy Cache Commands

local-user-db

Syntax **local-user-db** *local-user-db-name* [**create**]
no local-user-db *local-user-db-name*

Context config>subscr-mgmt

Description This command enables the context to configure a local user database.

The **no** form of this command removes the local user database name from the configuration.

Parameters *local-user-db-name* — Specifies the name of a local user database up to 32 characters.

dhcp

Syntax **dhcp**

Context config>subscr-mgmt>loc-user-db

Description This command configures DHCP host parameters.

host

Syntax **host**

Context config>subscr-mgmt>loc-user-db

Description This command enables the context to configure DHCP host parameters.

match-radius-proxy-cache

Syntax **match-radius-proxy-cache**

Context config>subscr-mgmt>loc-user-db>ipoe>host

Description This command enables the context to configure match-radius-proxy-cache parameters.

fail-action

Syntax **fail-action** {**continue** | **drop**}
no fail-action

Context config>subscr-mgmt>loc-user-db>ipoe>host>match-radprox-cache

Description This command specifies the router's action when failed to find matched radius-proxy-server cache entry.

The **no** form of this command reverts to the default.

Default fail-action drop

Parameters **continue** — Specifies that the will proceed with ESM authentication without dropping the DHCP packet.

drop — Specifies that the router will drop the DHCP packet.

mac-format

Syntax **mac-format** *format*
no mac-format

Context config>subscr-mgmt>loc-user-db>ipoe>host>match-radprox-cache

Description This command specifies the format of MAC address used for matching incoming DHCP DISCOVER against the RADIUS proxy cache.

The **no** form of this command reverts to the default.

Default mac-format "aa:"

Parameters *format* — Specifies the format string that specifies the format of MAC address.

Values

mac-format: (only when match is equal to mac)
like ab: for 00:0c:f1:99:85:b8
or XY- for 00-0C-F1-99-85-B8
or mmmm. for 0002.03aa.abff
or xx for 000cf19985b8

match

Syntax **match** {**circuit-id** | **mac** | **remote-id**}
match option [*option*] [**option6** [*option6*]]
match option6 [*option6*]
no match

Context config>subscr-mgmt>loc-user-db>ipoe>host>match-radprox-cache

Description This command specifies the field/option of DHCP packet that is used to match against the radius-proxy-server cache.

The **no** form of this command reverts to the default.

Default match mac

Parameters **circuit-id** — Specifies to match the circuit-id in DHCP option82.

remote-id — Specifies to match the remote
-id in DHCP option82.

mac — Specifies to match the MAC address of DHCP client

option — Specifies to use the specified DHCP option.

Values 1 to 254

option6 — Specifies to use the specified DHCP option.

Values 1 to 65535

number — Specifies the DHCPv6 option to retrieve the value to be used as lookup key in the RADIUS Proxy cache

Values 1 to 65535

server

Syntax **server** [**service** *service-id*] **name** *server-name*
no server

Context config>subscr-mgmt>loc-user-db>ipoe>host>match-radprox-cache

Description This command specifies the name of radius-proxy-server and optionally id of the service that the radius-proxy-server resides in.

The **no** form of this command removes the parameters from the configuration.

Parameters **service** *service-id* — Specifies the ID or name of the service.

Values 1 to 214748365
svc-name up to 64 char maximum

name *server-name* — Specifies the name of radius-proxy-server up to 32 characters.

11.25.2.5 WLAN-GW-Group Commands

wlan-gw-group

Syntax **wlan-gw-group** *group-id* [**create**] [**redundancy** *unit*]
no wlan-gw-group *group-id*

Context config>isa

Description This command creates a WLAN GW group that contains a set of ISAs to be used in WLAN-GW functionality. A WLAN-GW group can also be used where a NAT group is expected. The WLAN-GW group ID shares the same number space with the NAT group.

At most, one WLAN-GW group may be configured.

The optional redundancy parameter determines the provisioning and redundancy mode.

- IOM mode

  A whole IOM is added to the group. The IOM must be fully provisioned with BB ISA modules. In IOM mode, when a single ISA fails, the entire IOM is considered to have failed and all subscribers are recovered on a backup IOM.
- ISA mode

  BB ISA modules are added separately with no restriction put on other MDAs in the IOM. When a single ISA fails, a backup ISA will try to recover as many subscribers as possible but may run out of resources (for example, queues, policers, host entries) during the recovery process. It is recommended to pair ISAs with MDAs and services that do not consume many IOM resources.

The **no** form of this command removes the group.

Parameters *group-id* — Specifies WLAN Gateway Integrated Service Adaptor (ISA) Groups.

Values 1 to 4

unit — Specifies the provisioning and redundancy mode.

Values mda or iom

active-iom-limit

Syntax **active-iom-limit** *number*
no active-iom-limit

Context config>isa>wlan-gw-group

Description This command specifies the number of WLAN-GW IOMs used as active IOMs from the total number of configured WLAN-GW IOMs. If there are more configured IOM than active-iom-limit, then the remaining number of IOMs is designated as backup(s).

The **no** form of this command removes the number from the configuration.

Parameters *number* — Specifies the number of IOMs in this WLAN Gateway ISA group that are intended for active use.

Values 1 to 3

active-mda-limit

Syntax **active-mda-limit** *number*
no active-mda-limit

Context config>isa>wlan-gw-group

Description This command specifies how many ISAs may be in active use by the WLAN-GW group at the same time. If the maximum number of active ISAs is reached and more ISAs are added to the group, the new ISAs are considered to be in standby mode.

The **no** form of this command removes the limit on the maximum number of active ISAs.

Parameters *number* — Specifies the number of WLAN-GW ISAs intended for active use.

Values 1 to 14

distributed-sub-mgmt

Syntax [**no**] **distributed-sub-mgmt**

Context config>isa>wlan-gw-group

Description This command configures the WLAN gateway distributed subscriber management.

isa-aa-group

Syntax **isa-aa-group** *aa-group-id*
no isa-aa-group

Context config>isa>wlan-gw-group>distributed-sub-mgmt

Description This command configures an ISA application assurance group for WLAN gateway DSM subscribers.

isa-aa-oversubscription-factor

Syntax **isa-aa-oversubscription-factor** *factor*
no isa-aa-oversubscription-factor

Context config>isa>wlan-gw-group>distributed-sub-mgmt

Description This command specifies by how much an AA ISA is oversubscribed when linked to a WLAN-GW group. A factor of 1 indicates that each AA ISA is linked to a single WLAN-GW ISA, while a factor of 10 indicates that each AA ISA is linked to up to 10 WLAN-GW ISAs. The factor must be an integer but poses an oversubscription limit, not an exact ratio. For example, for 2 AA ISAs and 5 WLAN-GW ISAs, a factor of 3 or higher is valid. Additional standby ISAs can be added until the oversubscription limit is reached.

The **no** version of this command resets the configuration to the default value.

Default isa-aa-oversubscription-factor 1

Parameters *factor* — The number of WLAN GW ISAs that can be served by a single AA ISA.

Values 1 to 10

iom

Syntax **iom** *slot-number* **type** {[**load-balancer**] [**ue-anchor**]}
no iom *slot-number*

Context config>isa>wlan-gw-group

Description This command designates the specified IOM as a WLAN-GW IOM. Each WLAN-GW IOM must be provisioned with two ISA-BB modules on a hardware chassis and with an ISA-BB module in the first MDA slot in the VSR.

The **no** form of this command removes the IOM from the configuration.

Parameters *slot-number* — Indicates the IOM slot to be used in the WLAN-GW group.

Values 1 to 10

type {[**load-balancer**] [**ue-anchor**]} — This parameter is supported on the VSR only. It determines if an IOM slot is used for load-balancing or UE anchoring and processing, or both. When the **wlan-gw-group** has only a single IOM, it is required to put this IOM in both modes at the same time.

mda

Syntax [**no**] **mda** *mda-id*

Context config>isa>wlan-gw-group

Description This command enables an ISA for WLAN-GW functionality.

The **no** form of this command removes the ISA from the WLAN-GW configuration.

Parameters *mda-id* — Indicates the IOM and MDA slot in format *slot*/*mda*.

Values slot — 1 to 10

mda — 1 to 2

nat

Syntax **nat**

Context config>isa>wlan-gw-group

Description This command enables the context to configure NAT parameters under wlan-gw-group.

radius-accounting-policy

Syntax **radius-accounting-policy** *nat-accounting-policy*
no radius-accounting-policy

Context config>isa>wlan-gw-group>nat

Description This command configures the RADIUS accounting policy to use for each MDA in this ISA group.

The **no** form of this command removes the accounting policy from the configuration.

Parameters *nat-accounting-policy* — Specifies the RADIUS accounting policy up to 32 characters.

session-limits

Syntax **session-limits**

Context config>isa>wlan-gw-group>nat

Description This command configures the ISA NAT group session limits.

reserved

Syntax **reserved** *num-sessions*
no reserved

Context config>isa>nat>session-limits

Description This command configures the number of sessions per block that is reserved for prioritized sessions.

The **no** form of this command reverts to the default.

Parameters *num-sessions* — Specifies the number of sessions reserved for prioritized sessions.

Values 0 to 4194303

watermarks

Syntax **watermarks high** *percentage* **low** *percentage*
no watermarks

Context config>isa>nat>session-limits

Description This command configures the ISA NAT group watermarks.

The **no** form of this command reverts to the default.

Parameters *percentage* — Specifies the high watermark of the number of sessions for each MDA in this NAT ISA group.

Values 2 to 100

percentage — Specifies the low watermark of the number of sessions for each MDA in this NAT ISA group.

Values 1 to 99

suppress-lsn-events

Syntax [**no**] **suppress-lsn-events**

Context configure>isa>wlan-gw-group>nat

Description This command suppresses the generation of Large Scale NAT (LSN) events when RADIUS accounting is enabled.

By default, only one logging facility for tracking subscribers in LSN44, DS-lite, and NAT64 can be enabled at the time, either the SR OS event logging facility or the RADIUS logging facility. Note that SR OS event logs can be sent to multiple destinations, such as the console session, a telnet or SSH session, memory logs, file destinations, SNMP trap groups, and syslog destinations.

If RADIUS logging is enabled, the NAT logs are sent to the RADIUS destination and the NAT logs are suppressed in the SR OS event logging facility, for example, NAT logs are not sent to the syslog server.

If RADIUS logging is disabled, the NAT logs are sent to the SR OS event logging facility, for example, syslog, assuming that the events are enabled via the SR OS event-control (**config>log>event-control nat event generate**).

The **no** form of this command, the NAT logs can be sent to both logging facilities simultaneously, the SR OS event logging facility and RADIUS logging facility.

Default suppress-lsn-events

suppress-lsn-sub-blks-free

Syntax [**no**] **suppress-lsn-sub-blks-free**

Context configure>isa>wlan-gw-group>nat

Description This command suppresses the tmnxNatLsnSubBlksFree summary notification and use the tmnxNatPlBlockAllocationLsn notifications. When the SR OS node is in a state of excessive logging, the queue associated with the transmission of logs on the MS-ISA can become congested. This event further delays the generation of logs, and with this, further allocations and deallocations of NAT resources (port-blocks) is stalled until the queue is relieved of congestion. For example, an excessive logging state in the system can be caused by issuing a command to clear a large number of NAT subscribers where a large number of resources (port-blocks) are released at once.

The **suppress-lsn-sub-blks-free** command enables the generation of individual logs carried in event-id 2012 for every released port block regardless of the state of the transmission queue (whether congested or not). If NAT subscribers have a large number of allocated port blocks (this could be hundreds of port blocks per subscriber), generating individual logs per port-block release contributes to the congestion.

To alleviate transmission queue congestion, this behavior can be changed by disabling this command (**no suppress-lsn-sub-blks-free**). This causes the suppression of logs related to the release of individual port blocks of a NAT subscriber when the transmission queue is congested. As a result, only a summarized release log via event-id 2021 for the subscriber is generated. The purpose of this new log is to inform the operator in a single message that all ports blocks for the subscriber are released. For example, the log message for LSN is “LSN subscriber all blocks freed”. The benefit of such summarization (or log aggregation) is to alleviate the congestion of the transmission queue and consequently accelerate resource releases. An effect is the decreased granularity of information.

If summarization is enabled (**no suppress-lsn-sub-blks-free**) while there is no logging congestion in the system, the port block releases continue to be logged individually via the event-id 2012 (assuming that this is enabled in the event control), except for the last port block of the subscriber. When the last port block is released, the log with event-id 2021 is generated indicating that all port blocks for the subscriber are now released without carrying the specific information about this last port block that is released.

watermarks

Syntax **watermarks**

Context config>isa>wlan-gw-group

Description This command enables the context to configure ISA watermark notifications.

mark

Syntax **mark** *entity* **high** *percentage-high* **low** *percentage-low*
no mark *entity*

Context config>isa>wlan-gw-group>watermarks

Description This command enables a watermark notification. If the watermark is set, it generates a notification when the corresponding resource consumption goes above the high percentage. No additional notifications are sent until resource consumption goes under the low watermark, upon which, a notification is sent indicating the high watermark is no longer hit.

The **no** form of this command disables the watermark notification.

Parameters *entity* — Specifies which watermark to set.

Values user-equipment | bridge-domain | radius-proxy-client

percentage-high — Specifies the high watermark in percentage of total resources available.

Values 1 to 100

percentage-low — Specifies the low watermark in percentage of total resources available.

Values 0 to 99

11.25.2.6 Port Policy Commands

port-policy

Syntax **port-policy** *port-policy-name* [**create**]
no port-policy *port-policy-name*

Context config

Description This command either creates a new port-policy with create parameter or enters the configuration context of an existing port-policy.

The **no** form of this command removes the port policy name from the configuration.

Parameters *port-policy-name* — Specifies the name of *port-policy* up to 32 characters.

create — Keyword used to create a *port-policy*.

egress-scheduler-policy

Syntax **egress-scheduler-policy** *port-sched-plcy*
egress-scheduler-policy

Context config>port-policy

Description This command specifies the port-scheduler-policy to use in the egress direction for the internal port connecting the WLAN-GW IOM to the MS-ISA.

The **no** form of this command removes the policy from the configuration.

Parameters *port-sched-plcy* — Specifies the name of the port-scheduler-policy up to 32 characters.

11.25.2.7 WLAN-GW Group Interface Commands

Note: The **wlan-gw** commands apply only to the 7750 SR platform.

group-interface

Syntax **group-interface** *ip-int-name* [**create**]
group-interface *ip-int-name* [**create**] **lns**
group-interface *ip-int-name* [**create**] **wlangw**
no group-interface *ip-int-name* [**create**]

Context config>service>ies>subscriber-interface
config>service>vprn>subscriber-interface

Description This command creates a group interface. This interface is designed for triple-play services where multiple SAPs are part of the same subnet. A group interface may contain one or more SAPs.

The **no** form of this command removes the group interface from the subscriber interface.

Parameters *ip-int-name* — Specifies the interface name of a group interface. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

lns — Specifies to use LNS.

wlangw — Specifies the group interface for wlan-gw.

wlan-gw

Syntax **wlan-gw**

Context config>service>ies>sub-if>group-interface
config>service>vprn>sub-if> group-interface

Description This command enables the context to configure wlan-gw parameters.

vlan-tag-ranges

Syntax **vlan-tag-ranges**

Context config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw

Description This command enables the context to configure vlan-to-retail-map parameters to map dot1Q tags to retail-service-id. The WIFI AP could insert a dot1Q tag in the Layer 2 frame within the GRE tunnel to indicate the retail service provider for the subscriber.

range

Syntax **range start** [*range*] **end** [*range*]
range default
no range start [*range*] **end** [*range*]

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges
config>service>ies>sub-if>grp-if>wlan-gw>ranges

Description This command creates or enters the context of specified VLAN range for configuration applicable to that range of VLANs.

Default range default

Parameters **start** — Specifies the start of the VLAN range.

Values 0 to 4096

end — Specifies the end of VLAN the range.

Values 0 to 4096

default — Configures defaults for the interface.

authenticate-on-dhcp

Syntax [**no**] **authenticate-on-dhcp**

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range

Description This command enables initial authentication (when there is no state for the UE on the ISA), to be triggered by DHCP DISCOVER or REQUEST. The default behavior s authentication based on first Layer 3 packet.

The **no** form of this command reverts to the default.

authentication

Syntax **authentication**

Context config>service>vprn>sub-if>grp-if>wlan-gw
config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range

Description This command enables the context to create configuration for authenticating a user from the WLAN-GW ISA.

authentication-policy

Syntax **authentication-policy** *policy-name*
no authentication-policy

Context config>service>vprn>sub-if>grp-if>wlan-gw>authentication
config>service>ies>sub-if>grp-if>wlan-gw>authentication
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>authentication
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>authentication

Description This command specifies authentication policy configured under aaa context for authenticating users on WLAN-GW ISA.

The **no** form of this command removes the policy-name from the configuration.

Parameters *policy-name* — Specifies the name of the authentication policy up to 32 characters.

hold-time

Syntax **hold-time** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no hold-time

Context config>service>vprn>sub-if>grp-if>wlan-gw>authentication
config>service>ies>sub-if>grp-if>wlan-gw>authentication
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>authentication

config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>authentication

Description This command configures the minimum time that a user is held down after a failed authentication attempt.

The **no** form of this command reverts to the default.

Default hold-time sec 5

Parameters *hours* — Specifies the minimum time that a user is held down in hours.

Values 1 to 1

minutes — Specifies the minimum time that a user is held down in minutes.

Values 1 to 59

seconds — Specifies the minimum time that a user is held down in seconds.

Values 0to 59

vlan-mismatch-timeout

Syntax **vlan-mismatch-timeout** *seconds*
no vlan-mismatch-timeout

Context config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>authentication
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>authentication

Description This command configures the timeout value for the RADIUS proxy cache if a packet is received with a non-matching VLAN tag. The new timeout value is the lesser of the **vlan-mismatch-timeout** value and the currently remaining proxy cache timeout value.

The **no** form of this command disables the timeout behavior. The cache timeout value will remain unchanged.

Parameters *seconds* — Specifies the timeout value for the RADIUS proxy cache, in seconds.

Values 5 to 60

distributed-sub-mgmt

Syntax **distributed-sub-mgmt**

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range

Description This command enables the context to configure distributed-sub-mgmt configuration per vlan-range. This also includes vlan-range default, which makes this configuration applicable to the **wlan-gw group-interface**.

accounting-policy

Syntax **accounting-policy** *policy-name*
no accounting-policy

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dsm
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dsm

Description This command specifies the **isa-radius-policy** used for accounting messages originated from the ISAs in the **wlan-gw** group. The policy can specify up to five accounting servers and configuration-specific to these accounting servers. It also specifies configuration specific to RADIUS client on ISAs and RADIUS attributes to be included in accounting messages.

Parameters *policy-name* — Specifies the name of the account policy up to 32 characters.

accounting-update-interval

Syntax **accounting-update-interval** [*interval*]
no accounting-update-interval

Context config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dsm
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dsm

Description This command enables the interim accounting and specifies the interim accounting interval.

Parameters *interval* — Specifies the interim accounting interval in seconds.

Values 5 to 259200

collect-aa-acct-stats

Syntax [**no**] **collect-aa-acct-stats**

Context config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dsm
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dsm

Description This command enables Application Assurance account statistics collection.

def-app-profile

Syntax **def-app-profile** *profile-name*
no def-app-profile

Context config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dsm
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dsm

Description This command configures the default application profile.

dsm-ip-filter

Syntax **dsm-ip-filter** *dsm-ip-filter-name*
no dsm-ip-filter

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dsm
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dsm

Description This command configures an IP filter that is distributed on ISA cards.

This command specifies the IP filter applied to all UEs corresponding to default vlan-range (such as a group-interface) or the specified vlan-range. The IP filter can be created in the **config>subscr-mgmt>isa-filter** context, and can contain up to 1024 match entries. The IP filter can be overridden per UE from RADIUS via access-accept or COA.

The **no** form of this command reverts to the default.

Parameters *dsm-ip-filter-name* — Specifies the identifier of the distributed-sub-mgmt IP filter.

egress-policer

Syntax **egress-policer** [*policer*-name]
no egress-policer

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dsm
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dsm

Description This command specifies the egress policer applied to all UEs corresponding to default vlan-range (such as, group-interface) or the specified vlan-range. The policer can be created in the **config>subscr-mgmt>isa-policer** context. The egress policer can be overridden per UE from RADIUS via access-accept or COA.

The **no** form of this command reverts to the default.

Parameters *policer-name* — Specifies the identifier of the distributed-sub-mgmt policer for egress traffic up to 256 characters.

ingress-policer

Syntax **ingress-policer** *policer-name*
no ingress-policer

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dsm
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dsm

Description .This command specifies the ingress policer applied to all UEs corresponding to default vlan-range (such as group-interface) or the specified vlan-range. The policer can be created in the **config>subscr-mgmt>isa-policer** context. The ingress policer can be overridden per UE from RADIUS via access-accept or COA.

The **no** form of this command reverts to the default.

Parameters *policer-name* — Specifies the identifier of the distributed-sub-mgmt policer for ingress traffic.

one-time-redirect

Syntax **one-time-redirect url** *rdr-url-string* **port** *port-num*
no one-time-redirect

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dsm
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dsm

Description This command enables one-time http-redirect to specified redirect URL for traffic matching the specified destination port.

The **no** form of this command reverts to the default.

Parameters **url** *rdr-url-string* — Specifies the HTTP web address that is sent to the user's browser.

port *port-num* — Specifies the destination port number as a decimal hex or binary.

Values 1 to 65535

default-retail-svc-id

Syntax **default-retail-svc-id** *service-id*
no default-retail-svc-id

Context config>service>ies>sub-if>grp-if>wlan-gw>vlan-tag-ranges
config>service>vprn>sub-if>grp-if>wlan-gw>vlan-tag-ranges

Description This command specifies the id of default retail service if there is no match found in VLAN to retail map configuration (specified by the **vlan** command). For DSM and migrant, this command is only applicable for non-NAT stacks.

Parameters *service-id* — Specifies the identifier of the retail service to be used by default of a value in the retail service map of this interface

Values 1 to 2147483650
svc-name: up to 64 characters

vlan

Syntax **vlan start** [*value*] **end** [*value*] **retail-svc-id** *service-id*
no vlan start [*value*] **end** [*value*]

Context config>service>ies>sub-if>grp-if>wlan-gw>retailer
config>service>vprn>sub-if>grp-if>wlan-gw>retailer

Description This command creates a mapping from a range of VLANs (appearing in the wlan-gw encapsulated Layer 2 frame) to a retail service ID.

The **no** form of this command removes the parameters from the configuration.

Parameters **start** — Specifies the start VLAN tag of this range.

Values 0 to 4095

end — Specifies the end VLAN tag of this range.

Values 0 to 4095

retail-svc-id *service-id* — Specifies the identifier of the retail service to be used by default of a value in the retail service map of this interface.

Values 1 to 2147483650
svc-name: up to 64 characters

wlan-gw-group

Syntax **wlan-gw-group** *group-id*
no wlan-gw-group

Context config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw

Description This command specifies the ID of the **wlan-gw-group** that the **wlan-gw** gateway binds to.

The **no** form of this command removes the value from the **wlan-gw** configuration.

Parameters *group-id* — Specifies the ISA WLAN-GW group.

Values 1 to 4

ip-mtu

Syntax **ip-mtu** *octets*
no ip-mtu

Context config>service>ies>subscriber-interface
config>service>vprn>subscriber-interface

Description This command specifies the maximum size of frames on this group-interface. Packets larger than this will get fragmented.

The **no** form of this command removes this functionality.

Parameters *octets.* — Specifies the largest frame size (in octets) that this interface can handle

Values 512 to 9000

sap-parameters

Syntax **sap-parameters**

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command enables the context to configure parameters that can be applied to automatically-generated internal SAPs.

anti-spoof

Syntax **anti-spoof** {**ip-mac** | **nh-mac**}
no anti-spoof

Context config>service>ies>sub-if>grp-if>sap-parameters
config>service>vprn>sub-if>grp-if>sap-parameters

Description This command configures the anti-spoof type of the SAP.

The type of anti-spoof filtering defines what information in the incoming packet is used to generate the criteria to lookup an entry in the anti-spoof filter table. The type parameter (**ip-mac** or **nh-mac**) defines the anti-spoof filter type enforced by the SAP when anti-spoof filtering is enabled.

The **no** form of this command reverts to the default.

Default anti-spoof ip-mac

Parameters **ip-mac** — Configures SAP anti-spoof filtering to use both the source IP address and the source MAC address in its lookup. The **anti-spoof ip-mac** command will fail if the default anti-spoof filter type of the SAP is ip-mac and the default is not overridden, or if the SAP does not support Ethernet encapsulation.

nh-mac — Indicates that the ingress anti-spoof is based on the source MAC address and egress anti-spoof is based on the nh-ip-address

sub-sla-mgmt

Syntax [**no**] **sub-sla-mgmt**

Context config>service>ies>sub-if>grp-if>sap-parameters
config>service>vprn>sub-if>grp-if>sap-parameters

Description This command enables the context to configure subscriber management parameters.

The **no** form of this command removes the parameters from the configuration.

Default sub-sla-mgmt

def-app-profile

Syntax **def-app-profile** *profile-name*
no def-app-profile

Context config>service>ies>sub-if>grp-if>sap-parameters>sub-sla-mgmt
config>service>vprn>sub-if>grp-if>sap-parameters>sub-sla-mgmt

Description This command configures the default application profile.

The **no** form of this command removes the profile name from the configuration.

Parameters *profile-name* — Specifies the default application profile name up to 32 characters

def-sla-profile

Syntax **def-sla-profile** *default-sla-profile-name*
no def-sla-profile

Context config>service>ies>sub-if>grp-if>sap-parameters>sub-sla-mgmt
config>service>vprn>sub-if>grp-if>sap-parameters>sub-sla-mgmt

Description This command specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sla-profile** context.

An SLA profile is a named group of QoS parameters used to define per service QoS for all subscriber hosts common to the same subscriber within a provider service offering. A single SLA profile may define the QoS parameters for multiple subscriber hosts. SLA profiles are maintained in two locations, the subscriber identification policy and the subscriber profile templates. After a subscriber host is associated with an SLA profile name, either the subscriber identification policy used to identify the subscriber or the subscriber profile associated with the subscriber host must contain an SLA profile with that name. If both the subscriber identification policy and the subscriber profile contain the SLA profile name, the SLA profile in the subscriber profile is used.

The **no** form of this command removes the default SLA profile from the SAP configuration.

Parameters *default-sla-profile-name* — Specifies a default SLA profile for this SAP. The SLA profile must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sla-profile** context.

def-sub-id

Syntax **def-sub-id string** *sub-id*
def-sub-id use-auto-id
no def-sub-id

Context config>service>vprn>sub-if>grp-if>sap-parameters>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap-parameters>sub-sla-mgmt

Description This command configures the default subscriber ID. The default is used if no other source (like RADIUS) provides a subscriber identification string.

Parameters *sub-id* — Specifies the default subscriber identification up to 32 characters

use-auto-id — Specifies that the auto-generated subscriber identification string, is used as the default subscriber identification string

def-sub-profile

Syntax **def-sub-profile** *sub-profile-name*

Context config>service>vprn>sub-if>grp-if>sap-parameters>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap-parameters>sub-sla-mgmt

Description This command specifies a default subscriber profile. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sub-profile** context.

A subscriber profile defines the aggregate QoS for all hosts within a subscriber context. This is done through the definition of the egress and ingress scheduler policies that govern the aggregate SLA for subscriber using the subscriber profile. Subscriber profiles also allow for specific SLA profile definitions when the default definitions from the subscriber identification policy must be overridden.

The **no** form of this command removes the default SLA profile from the configuration.

Parameters *sub-profile-name* — Specifies a default subscriber profile. The subscriber profile must be defined in the **config>subscr-mgmt>sub-profile** context.

sub-ident-policy

Syntax **sub-ident-policy** *sub-ident-policy-name*

Context config>service>vprn>sub-if>grp-if>sap-parameters>sub-sla-mgmt
config>service>ies>sub-if>grp-if>sap-parameters>sub-sla-mgmt

Description This command associates a subscriber identification policy. The subscriber identification policy must be defined in the **config>subscr-mgmt>sub-ident-policy** context.

Subscribers are managed by the system through the use of subscriber identification strings. A subscriber identification string uniquely identifies a subscriber. For static hosts, the subscriber identification string is explicitly defined with each static subscriber host.

For dynamic hosts, the subscriber identification string must be derived from the DHCP ACK message sent to the subscriber host. The default value for the string is the content of Option 82 CIRCUIT-ID and REMOTE-ID fields interpreted as an octet string. As an option, the DHCP ACK message may be processed by a subscriber identification policy which has the capability to parse the message into an alternative ASCII or octet string value.

When multiple hosts on the same port are associated with the same subscriber identification string they are considered to be host members of the same subscriber.

The **no** form of this command removes the default subscriber identification policy from the configuration.

Parameters *sub-ident-policy-name* — Specifies a subscriber identification policy for this SAP. The subscriber profile must be defined prior to associating the profile with a SAP in the **config>subscr-mgmt>sub-ident-policy** context.

egress

Syntax **egress**

Context config>service>ies>sub-if>grp-if>wlan-gw

Description This command enables the context to configure egress QoS parameters for wlan-gw tunnels.

rate

Syntax **rate** {**max** | **rate**}
no rate

Context config>service>ies>sub-if>grp-if>wlan-gw>egress>agg-rate
config>service>vprn>sub-if>grp-if>wlan-gw>egress>agg-rate

Description This command defines the enforced aggregate rate for all queues associated with the agg-rate context. A rate must be specified for the agg-rate context to be considered to be active on the context's object (SAP, subscriber, Vport, and so on).

The **no** form of this command reverts to the default.

agg-rate-limit

Syntax **agg-rate-limit** *kilobits-per-second*
no agg-rate-limit

Context config>service>ies>sub-if>grp-if>wlan-gw>egress
config>service>vprn>sub-if>grp-if>wlan-gw>egress

Description This command controls an HQoS aggregate rate limit. It is used in conjunction with the following parameter commands: **rate**, **limit-unused-bandwidth**, **and queue-frame-based-accounting**.

The **no** form of this command removes the rate from the configuration.

Parameters *kilobits-per-second* — Specifies the aggregate rate limit

Values 1 to 100000000, **max**

hold-time

Syntax **hold-time infinite**
hold-time [**time**]
no hold-time

Context config>service>ies>sub-if>grp-if>wlan-gw>egress
config>service>vprn>sub-if>grp-if>wlan-gw>egress

Description This command configures the time for which egress shaping resources associated with a wlan-gw tunnel are held after the last subscriber on a tunnel is deleted.

Parameters *time* — Specifies the time, in seconds, for which shaping resources are held in seconds after last subscriber is deleted

Values infinite to 1 to 86400

qos

Syntax **qos** *policy-id*
no qos

Context config>service>ies>sub-if>grp-if>wlan-gw>egress

Description This command configures the identifier of the egress QoS policy associated with each wlan-gw tunnel of this interface.

The **no** form of this command removes the policy ID from the configuration.

Default qos 1

Parameters *policy-id* — Specifies to apply the specified sap-*egress-policy-id*

Values 1 to 65535
name: A string up to 64 characters

scheduler-policy

Syntax **scheduler-policy** *scheduler-policy-name*
no scheduler-policy

Context config>service>ies>sub-if>grp-if>wlan-gw>egress

Description This command configures the identifier of the egress scheduler policy associated with each wlan-gw tunnel of this interface.

The **no** form of this command removes the scheduler policy name from the configuration.

Parameters *scheduler-policy-name* — Specifies the identifier of the egress scheduler policy associated with each wlan-gw tunnel of this interface

shape-multi-client-only

Syntax [**no**] **shape-multi-client-only**

Context config>service>ies>sub-if>grp-if>wlan-gw>egress

Description This command enables the egress shaping is only enabled for a wlan-gw tunnel while there are multiple UE (User Equipment) using it.

The **no** form of this command disables the egress shaping.

shaping

Syntax **shaping** {**per-retailer** | **per-tunnel**}
no shaping

Context config>service>ies>sub-if>grp-if>wlan-gw>egress

Description This command configures the granularity of the egress shaping for wlan-gw on this group interface.

The **no** form of this command removes the parameter from the configuration.

Parameters **per-tunnel** — Specifies that a separate shaper is applied to each wlan-gw tunnel

per-retailer — Specifies that a separate shaper is applied to each retailer Mobile Network Operator's fraction of the wlan-gw tunnel payload

gw-address

Syntax **gw-address** *ip-address*
no gw-address

Context config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw

Description This command specifies gateway endpoint address for the wlan-gw tunnel.

The **no** form of this command removes the value from the wlan-gw configuration.

Parameters *ip-address* — Specifies the IP address of the wlan-gw tunnels on this group interface

gw-ipv6-address

Syntax **gw-ipv6-address** *ipv6-address*
no gw-ipv6-address

Context config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw

Description This command specifies a gateway IPv6 endpoint address for the wlan-gw tunnel.

The **no** form of this command removes the IPv6 the gateway IPv6 endpoint address for the wlan-gw tunnel.

Parameters *ipv6-address* — Specifies the gateway IPv6 endpoint address

Values

ipv6-address :	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D

learn-ap-mac

Syntax **learn-ap-mac** [**delay-auth**]
no learn-ap-mac

Context config>service>vprn>sub-if>grp-if>wlan-gw
config>service>ies>sub-if>grp-if>wlan-gw

Description This command enables the sending of ARP or ND packets on the wlan-gw GRE tunnel upon certain events. The target IP address in the ARP/ND packet is the endpoint IP address of the AP. The ARP/ND response from the AP should contain the AP MAC, which subsequently can be reported in called-station-id. When enabled this is sent for following events:

- CPM: Mobility to an AP for which the AP-MAC is not yet known.
- CPM: RS-triggered authentication on an AP for which the AP-MAC is not yet known
- ISA: Any mobility event
- ISA: Any authentication where the AP-MAC is not yet known (for example, from RADIUS proxy cache, DHCP circuit-id, and so on). If the optional keyword **delay-auth** is provided the authentication is delayed until the ARP/ND is answered or timed out, after which the AP-MAC can be included in authentication.

This configuration is ignored for l2-ap and l2tpv3 access.

Parameters **delay-auth** — Specifies that authentication is delayed until the ARP/ND is answered or timed out, after which the AP-MAC can be included in authentication

l2-access-points

Syntax **l2-access-points**

Context config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw

Description This command enables the context to configure Layer 2 Access Points in WLAN Gateway Group-Interfaces.

l2-ap

Syntax **l2-ap** *sap-id* [**create**]
no l2-ap *sap-id*

Context config>service>vprn>sub-if>grp-if>wlan-gw>l2-access-points
config>service>ies >sub-if>grp-if>wlan-gw>l2-access-points

Description This command adds a specific SAP where Layer-2 WLAN-GW aggregation is performed. The following SAPs are supported.

- Ethernet
- LAG
- MPLS pseudowire SDPs

This command can be repeated multiple times to create multiple Layer-2 access points.

The **no** form of this command removes the Layer-2 access point. This is only allowed if the l2-ap SAP is shutdown.

Parameters *sap-id* — Specifies SAP to be created.

create — Keyword used to create the Layer-2 WLAN-GW aggregation instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

encap-type

Syntax **encap-type** {**default** | **null** | **dot1q** | **qinq**}
no encap-type

Context config>service>vprn>sub-if>grp-if>wlan-gw>l2-access-points>l2-ap
config>service>ies >sub-if>grp-if>wlan-gw>l2-access-points>l2-ap

Description If different from default, this command overrides the value specified by **l2-ap-encap-type** on wlan-gw level. See the description of l2-ap-encap-type for more detail. This value can only be changed while the l2-ap is shutdown.

The **no** form of this command sets the default value.

Default encp-type default

Parameters **default** — Specifies to use the value specified by l2-ap-encap-type

null — Specifies to use both the SAP and the AP are not VLAN-tagged

dot1q — Specifies to use either the AP or the SAP uses one VLAN tag

qinq — Up to two VLAN tags are used by the AP or SAP

epipe-sap-template

Syntax **epipe-sap-template** *name*
no epipe-sap-template

Context config>service>vprn>sub-if>grp-if>wlan-gw>l2-access-points>l2-ap
config>service>ies >sub-if>grp-if>wlan-gw>l2-access-points>l2-ap

Description This command specifies which SAP parameter template should be applied to the l2-ap SAP. This can only be changed when the l2-ap is shutdown.

The **no** form of this command removes the template, the SAP will use default parameters.

Parameters *name* — Specifies the name of the template to use

shutdown

Syntax **shutdown** *sap-id* [**create**]
no shutdown *sap-id*

Context config>service>vprn>sub-if>grp-if>wlan-gw>l2-access-points>l2-ap
config>service>ies >sub-if>grp-if>wlan-gw>l2-access-points>l2-ap

Description This command administratively enables this SAP to begin accepting Layer 2 packets for WIFI offloading.

The **no** form of this command disables this SAP.

Default shutdown

l2-ap-auto-sub-id-fmt

Syntax **l2-ap-auto-sub-id-fmt** {**include-ap-tags** | **sap-only**}

Context config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw

Description This command configures the contents of the auto-generated subscriber ID when the **ipoe-sub-id-key** command is set to include **sap-id** and the **def-sub-id** command is configured with **use-auto-id**. The VLANs must be configured so that the subscriber ID length is not exceeded.

This command can include either the SAP or the SAP + AP delimiting tags.

The **no** form of this command reverts to the default configuration.

Default l2-ap-auto-sub-id-fmt include-ap-tags

Parameters **include-ap-tags** — Specifies that the SAP + AP delimiting tags is used.

sap-only — Specifies that the SAP only is used.

l2-ap-encap-type

Syntax **l2-ap-encap-type** {**null** | **dot1q** | **qinq**}

Context config>service>vprn>sub-if>grp-if>wlan-gw
config>service>ies >sub-if>grp-if>wlan-gw

Description This parameter specifies the number of AP identifying VLAN tags for an AP. This is the default value that can be overridden per SAP. This value should at least be equal to the number of VLANs configured in the SAP or enabling a SAP will fail.

A SAP VLAN is explicitly configured, for example **l2-ap 1/1/1:25**. Other VLANs on the same port can still be used in other contexts.

The number of VLAN tags Epiped to WLAN-GW IOM equal the **l2-ap-encap-type** minus the encaps of the SAP. Upon receipt of a packet these VLANs is stored as a Layer 2 tunnel identifier, and are only used in context of WLAN-GW.

The **no** form of this command sets the default value.

Default l2-ap-encap-type null

Parameters **null** — Both the SAP and the AP are not VLAN-tagged

dot1q — Either the AP or the SAP uses one VLAN tag

qinq — Up to two VLAN tags are used by the AP or SAP

mobility

Syntax **mobility**

Context config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw

Description This command enables the context to configure mobility parameters.

hold-time

Syntax **hold-time** *time*
no hold-time

Context config>service>ies>sub-if>grp-if>wlan-gw>mobility
config>service>vprn>sub-if>grp-if>wlan-gw>mobility

Description This command configures the minimum time that a User Equipment is held associated with its current Access Point (AP) before being associated with a new AP.

The hold time is used to prevent overwhelming the system with mobility triggers, by limiting the rate at which a UE can move from one AP to another while the system is very busy already.

Default hold-time 5

Parameters *time* — Specifies a hold-down time, in seconds, for handling of successive mobility triggers for a UE. It is the minimal time a UE stays associated with an AP.

Values 0 to 255

inter-vlan

Syntax [**no**] **inter-vlan**

Context config>service>ies>sub-if>grp-if>wlan-gw>mobility
config>service>vprn>sub-if>grp-if>wlan-gw>mobility

Description This command enables mobility within different VLANs of the same range. When enabled, mobility between different VLANs in a single vlan-range is allowed for the configured mobility triggers.

The **no** form of this command disables mobility between VLANs.

trigger

Syntax **trigger** [**data**] [**iapp**] [**control**]
no trigger

Context config>service>ies>sub-if>grp-if>wlan-gw>mobility
config>service>vprn>sub-if>grp-if>wlan-gw>mobility

Description This command specifies the type of packet used as a mobility trigger.

The **no** form of this command removes the parameters from the configuration and disables data-plane mobility.

Parameters **data** — Specifies that data traffic be used as a trigger

iapp — Specifies that Inter Access Point Protocol (IAPP) messages be used as a trigger

control — Specifies that control traffic can be used as a trigger

multi-tunnel-type

Syntax [**no**] **multi-tunnel-type**

Context config>service>ies>sub-if>grp-if>wlan-gw>mobility
config>service>vprn>sub-if>grp-if>wlan-gw>mobility

Description This command enables terminating multiple types of tunnels.

The **no** form of this command disables terminating multiple types of tunnels.

no multi-tunnel-type

oper-down-on-group-degrade

Syntax [**no**] **oper-down-on-group-degrade**

Context config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw

Description This command operationally brings down the WLAN-GW group if the total number of operational WLAN-GW IOMs in the WLAN-GW group fall below the configured number of active WLAN-GW IOMs. This triggers withdrawal of the route to tunnel endpoint and subscriber subnets in routing.

router

Syntax **router** *router-instance*
no router

Context config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw

Description This command specifies the routing instance that wlan-gw gateway endpoint resides in.

The **no** form of this command removes the value from the wlan-gw configuration.

Default router

Parameters *router-instance* — Specifies the identifier of the virtual router instance where the tunneled User Equipment traffic is routed.

tcp-mss-adjust

Syntax **tcp-mss-adjust** *segment-size*
no tcp-mss-adjust

Context config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw

Description This command configures the TCP Maximum Segment Size (MSS) adjustment for the wlan-gw gateway.

The **no** form of this command disables adjusting tcp-mss values.

For DSM, this only applies to packets sent in the downstream direction (TCP SYN towards UE). For the upstream direction, it is also required to configure MSS adjust under the applicable NAT-policy.

Parameters *segment-size* — Specifies the value to put into the TCP Maximum Segment Size (MSS) option if not already present, or if the present value is higher.

Values 160 to 10240

tunnel-encaps

Syntax **tunnel-encaps**

Context config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw

Description This command enables the context to configure tunnel encapsulation parameters.

learn-l2tp-cookie

Syntax **learn-l2tp-cookie** {**if-match** | **never** | **always**} [**cookie** *hex string*]
no learn-l2tp-cookie

Context config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw

Description This command specifies when this system will learn the cookie from L2TP tunnels terminating on this interface. Learning the cookie means that the value of the octets 3-8 of the cookie is interpreted as an access point's MAC address, and used as such, for example in the Called-Station-Id attribute of RADIUS Interim-Update messages.

Parameters **if-match** — Specifies that the cookie is interpreted only if the value of the first two octets of the cookie is equal to the value of the object tmnxWlanGwSoftGreIfL2tpCookie.

cookie *hex string* — Specifies the value used to compare the first two bytes of the cookie. This parameter is only valid if **if-match** is configured.

Values 0x0000 to 0xFFFF...(4 hex nibbles)

never — Specifies that the cookie value will always be ignored.

always — Always learn the AP-MAC from the cookie, regardless of the value of the first two bytes.

retail-svc-id

Syntax **retail-svc-id** *service-id*
no retail-svc-id

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range

Description This command configures the retailer service.

Parameters *service-id* — Specifies the identifier of the retail service.

Values 1 to 2147483650
svc-name: up to 64 characters

router-advertisements

Syntax [**no**] **router-advertisements**

Context config>service>vprn>sub-if>grp-if>ipv6
config>service>ies>sub-if>grp-if>ipv6

Description This command configures IPv6 router advertisements for this group-interface.

current-hop-limit

Syntax [**no**] **current-hop-limit** *limit*

Context config>service>vprn>sub-if>grp-if>ipv6>rtr-adv
config>service>ies>sub-if>ipv6>rtr-adv

Description This command configures the hop-limit advertised for this group-interface.

Default current-hop-limit

Parameters *limit* — Specifies the default value to be placed in the current hop limit field in router advertisements sent from this interface

Values 0 to 255

pool-manager

Syntax **pool-manager**

Context config>service>ies>sub-if>wlan-gw
config>service>vprn>sub-if>wlan-gw

Description This command enables the context to configure pool manager data for a WLAN GW subscriber interface.

dhcpv6-client

Syntax **dhcpv6-client**

Context config>service>ies>sub-if>wlan-gw>pool-manager
config>service>vprn>sub-if>wlan-gw>pool-manager

Description This command configures the DHCPv6 client for the pool manager.

dhcpv4-nat

Syntax **dhcpv4-nat**

Context config>service>ies>sub-if>wlan-gw>pool-manager>dhcp6-client
config>service>vprn>sub-if>wlan-gw>pool-manager>dhcp6-client

Description This node enables address pools for DHCPv4 NAT inside addresses. This configuration is only available in wholesale interfaces.

ia-na

Syntax **ia-na**

Context config>service>ies>sub-if>wlan-gw>pool-manager>dhcp6-client
config>service>vprn>sub-if>wlan-gw>pool-manager>dhcp6-client

Description This command configures the IA-NA for the DHCPv6 client.

link-addr

Syntax **link-addr** *ipv6-address*
no link-addr

Context config>service>ies>sub-if>wlan-gw>pool-manager>dhcp6-client>slaac
config>service>vprn>sub-if>wlan-gw>pool-manager>dhcp6-client>slaac
config>service>ies>sub-if>wlan-gw>pool-manager>dhcp6-client>ia-na
config>service>vprn>sub-if>wlan-gw>pool-manager>dhcp6-client>ia-na
config>service>ies>sub-if>wlan-gw>pool-manager>dhcp6-client>dhcpv4-nat
config>service>vprn>sub-if>wlan-gw>pool-manager>dhcp6-client>dhcpv4-nat

Description This command specifies the ipv6-address that should be included in the link-address field of the relay header. This can be used for pool selection by the DHCPv6 server.

The **no** form of this command falls back to the default.

Parameters *ipv6-address* — Specifies the IPv6 address up to 32 characters.

pool-name

Syntax **pool-name** *name*
no pool-name

Context config>service>ies>sub-if>wlan-gw>pool-manager>dhcp6-client>slaac
config>service>vprn>sub-if>wlan-gw>pool-manager>dhcp6-client>slaac
config>service>ies >sub-if>wlan-gw>pool-manager>dhcp6-client>ia-na

config>service>vprn>sub-if>wlan-gw>pool-manager>dhcp6-client>ia-na
config>service>ies>sub-if>wlan-gw>pool-manager>dhcp6-client>dhcpv4-nat
config>service>vprn>sub-if>wlan-gw>pool-manager>dhcp6-client>dhcpv4-nat

Description This command specifies the pool name that should be sent in the DHCPv6 messages. This is reflected in the Nokia vendor specific pool option (vendor-id 6527, option-id 0x02).

The **no** form of this command removes pool-name and the option will not be sent in DHCPv6.

Parameters *name* — Specifies the pool name up with 32 characters,

lease-query

Syntax **lease-query** [**max-retry** *Max nbr of retries*]
no lease-query

Context config>service>vprn>sub-if>wlan-gw>pool-manager>dhcp6-client
config>service>ies>sub-if>wlan-gw>pool-manager>dhcp6-client

Description This command enables lease-query. If this is specified the dhcp6-client will retrieve any existing addresses when becoming active. The lease-query is performed for all of the configured servers

The **no** form of this command disables lease-query.

Parameters *Max nbr of retries* — Specifies the maximum number of retries before the lease query assumes no existing subnets were allocated.

Values 0 to 10

server

Syntax **server** *ipv6-address* [*ipv6-address*]
no server [*ipv6-address* [*ipv6-address*]]

Context config>service>vprn>sub-if>wlan-gw>pool-manager>dhcp6-client
config>service>ies>sub-if>wlan-gw>pool-manager>dhcp6-client

Description This specifies the DHCPv6 servers that are used for requesting addresses. Up to 8 servers can be used simultaneously.

The **no** form of this command removes the server. This cannot be executed while any DHCPv6 client application is not shutdown.

Parameters *ipv6-address* — Specifies the unicast IPv6 address of a dhcp6 server

slaac

Syntax **slaac**

Context config>service>vprn>sub-if>wlan-gw>pool-manager>dhcp6-client
config>service>ies>sub-if>wlan-gw>pool-manager>dhcp6-client

Description This command configures SLAAC for the DHCPv6 client.

source-ip

Syntax **source-ip** *ipv6-address*
no source-ip

Context config>service>vprn>sub-if>wlan-gw>pool-manager>dhcp6-client
config>service>ies>sub-if>wlan-gw>pool-manager>dhcp6-client

Description This command specifies the source-ip to be used by the DHCPv6 client.

The **no** form of this command removes the specific source-ip. In this case the DHCPv6 client will fall back to the IP address configured on the outgoing interface.

Parameters *ipv6-address* — Specifies the IPv6 address up to 32 characters

watermarks

Syntax **watermarks high** *high-percentage* **low** *low-percentage*
no watermarks

Context config>service>vprn>sub-if>wlan-gw>pool-manager
config>service>ies>sub-if>wlan-gw>pool-manager

Description This command configures the watermarks used to determine if a new prefix should be allocated or an old prefix should be removed. A new prefix is allocated when the total usage level for the ISA reaches the high watermark. A prefix is freed if no addresses are currently in use and the usage level without this prefix would be below the low watermark.

The **no** form of this command resets the watermarks to its default values of 95% high and 90% low.

Default watermarks high 95 low 90

Parameters *high-percentage* — Specifies the high watermark.

Values 80 to 99

low-percentage — Specifies the low watermark. The value must be lower than the high percentage value.

Values 50 to 98

wlan-gw-group

Syntax **wlan-gw-group** *nat-group-id*
no wlan-gw-group

Context config>service>vprn>sub-if>wlan-gw>pool-manager
config>service>ies>sub-if>wlan-gw>pool-manager

Description This command specifies the ISA WLAN gateway group.

Parameters *nat-group-id* — Specifies the identifier of the WLAN gateway group.

Values 1 to 4

redundancy

Syntax **redundancy**

Context config>service>ies>sub-if>wlan-gw

Description This command enables the context to configure WLAN-GW redundancy-related parameters.

export

Syntax **export** *ip-prefix/length*
no export

Context config>service>ies>sub-if>wlan-gw>redundancy

Description This command specifies an IPv4 route (prefix/length) per subscriber-interface to be exported (announced) to indicate liveness of the subscriber-interface on the WLAN-GW. This route is the one that is monitored in routing by the peer WLAN-GW to decide its state with respect.

The no form of this command reverts to the default.

Parameters *ip-prefix/length* — Specifies the IP prefix and length

Values ip-prefix:a.b.c.d
ip-prefix-length: 0 to 32

monitor

Syntax **monitor** *ip-prefix/length*
no monitor

Context config>service>ies>sub-if>wlan-gw>redundancy

Description This command specifies an IPv4 route (prefix/length) per subscriber-interface to be monitored in the FDB to determine liveness of the subscriber-interface (and consequently all associated group-interfaces of type wlangw) on a peer WLAN-GW. This route is the one that is advertised in routing by the peer WLAN-GW when the subscriber-interface and WLAN-GW group are operationally up

Parameters *ip-prefix/length* — Specifies the IP prefix and length

Values ip-prefix:a.b.c.d
ip-prefix-length: 0 to 32

11.25.2.8 Migrant User Support Commands

http-redirect-policy

Syntax **http-redirect-policy** *policy-name*
no http-redirect-policy

Context config>subscr-mgmt

Description This command configures the redirect policy to constrain forwarding of an unauthenticated “migrant” WIFI user.

Parameters *policy-name* — Specifies the HTTP redirect policy name up to 32 characters.

forward-entries

Syntax **forward-entries**

Context config>subscr-mgmt>http-rdr-plcy

Description Enters the context to configure entries that need to be forwarded

dst-port

Syntax **dst-port** *tcp-port*
no dst-port

Context config>subscr-mgmt>http-rdr-plcy

Description This command specifies the port to match the destination port in the HTTP request.

HTTP traffic that does not match this port, is not redirected.

The **no** form of this command reverts to the default.

Default dst-port 80

Parameters *tcp-port* — Specifies the TCP port.

Values 1 to 65535

dst-ip

Syntax **dst-ip** *ip-address* **protocol** *ip-protocol* **dst-port** *port-number*
dst-ip *ip-address* **protocol** *ip-protocol* **dst-port** *port-number* **prefix-length** *prefix-length*
no dst-ip *ip-address* **protocol** *ip-protocol* **dst-port** *port-number*

Context config>subscr-mgmt>http-rdr-plcy

Description This command configures traffic flow to be forwarded via match in the redirect policy.

Parameters *ip-address* — Specifies the IPv4 or IPv6 address to match the destination address in the IP header of the traffic received from the subscriber.

prefix-length — Specifies the length of the prefix specified by the ip-address.

Values 1 to 128 for IPv6
1 to 32 for IPv4

ip-protocol — Specifies the protocol to match the IP protocol in the IP header of the traffic received from the subscriber.

Values tcp, udp

port-number — Specifies the port to match the destination port in the HTTP request.

Values 1 to 65535

portal-hold-time

Syntax **portal-hold-time** *seconds*
no portal-hold-time

Context config>subscr-mgmt>http-rdr-plcy

Description This command configures the time for which the forwarding state applicable during redirect phase is held in the system, after the user has been authenticated on the portal. This allows the HTTP response from the portal to be forwarded back on the existing connection.

Parameters *seconds* — Specifies how long the system holds on to re-direct forwarding resources of a subscriber, after it has left the re-direct portal.

Values 1 to 60

url

Syntax **url** *rdr-url-sting*
no url

Context config>subscr-mgmt>http-rdr-plcy

Description This command configures the HTTP URL to re-direct the matching traffic to. It also can specify inclusion of original URL, MAC address and IP address of the subscriber in the redirect URL.

Parameters *rdr-url-sting* — Specifies the URL to redirect to

Values

rdr-url-string	[255 chars max]
macro substitutions:	
$URL	Request-URI in the HTTP GET Request received
$MAC	A string that represents the MAC address of the subscriber host
$IP	A string that represents the IP address of the subscriber host

wlan-gw

Syntax **wlan-gw**

Context config>service>vprn>sub-if>grp-if
config>service>ies>sub-if>grp-if

Description This command enables the context to configure wlan-gw parameters.

vlan-tag-ranges

Syntax **vlan-tag-ranges**

Context config>service>vprn>sub-if>grp-if>wlan-gw
config>service>ies>sub-if>grp-if>wlan-gw

Description This command enables the context for per vlan range configuration.

default-retail-svc-id

Syntax **default-retail-svc-id** *service-id*

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges
config>service>ies>sub-if>grp-if>wlan-gw>ranges

Description This command configures the default retailer service for WIFI users.

dhcp

Syntax **dhcp**

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges
config>service>ies>sub-if>grp-if>wlan-gw>ranges

Description This command enables the context to create DHCP configuration for WLAN-GW ISA subscribers (such as migrant subscribers).

dhcp6

Syntax **dhcp6**

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges
config>service>ies>sub-if>grp-if>wlan-gw>ranges

Description This command enables the context to create DHCP6 configuration for WLAN-GW ISA subscribers.

active-preferred-lifetime

Syntax **active-preferred-lifetime** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no active-preferred-lifetime

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dhcp6
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dhcp6
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>slaac
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>slaac

Description This command specifies the signaled preferred lifetime in DHCPv6 or SLAAC after full authentication. This is only applicable to DSM.

The **no** form of this command reverts to the default.

Default active-preferred-lifetime min 10

Parameters *hours* — Specifies the number of active preferred lifetime hours.

Values 1 to 1

minutes — Specifies the number of active preferred lifetime minutes.

Values 5 to 59

seconds — Specifies the number of active preferred lifetime seconds.

Values 1 to 59

active-valid-lifetime

Syntax **active-valid-lifetime** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no active-valid-lifetime

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dhcp6
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dhcp6
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>slaac
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>slaac

Description This command specifies the signaled valid lifetime in DHCPv6 or SLAAC after full authentication. This is only applicable to DSM.

The **no** form of this command reverts to the default.

Default active-valid-lifetime min 10

Parameters *hours* — Specifies the number of active-valid-lifetime hours.

Values 1 to 1

minutes — Specifies the number of active-valid-lifetime minutes.

Values 5 to 59

seconds — Specifies the number of active-valid-lifetime seconds.

Values 1 to 59

active-lease-time

Syntax **active-lease-time** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no active-lease-time

Context config>service>vprn>sub-if>grp-if>wlan-gw>dhcp
config>service>ies>sub-if>grp-if>wlan-gw>dhcp
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dhcp
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dhcp

Description This command configures the lease time for an authenticated user.

Default active-lease-time min 10

Parameters *hours* — Specifies the number of active lease time hours.

Values 1 to 1

minutes — Specifies the number of active lease time minutes.

Values 5 to 59

seconds — Specifies the number of active lease time seconds.

Values 1 to 59

initial-preferred-lifetime

Syntax **initial-preferred-lifetime** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no initial-preferred-lifetime

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dhcp6
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dhcp6
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>slaac
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>slaac

Description This command specifies the signaled preferred lifetime in DHCPv6 or SLAAC after full authentication (DSM and/or ESM).

The **no** form of this command reverts to the default.

Default initial-preferred-lifetime min 5

Parameters *hours* — Specifies the number of initial preferred lifetime hours.

Values 1 to 1

minutes — Specifies the number of initial preferred lifetime minutes.

Values 5 to 59

seconds — Specifies the number of initial preferred lifetime seconds.

Values 1 to 59

Combined values: min 5 – hrs 1

initial-valid-lifetime

Syntax **initial-valid-lifetime** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no initial-valid-lifetime

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dhcp6
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dhcp6
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>slaac
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>slaac

Description This command specifies the signaled preferred lifetime in DHCPv6 or SLAAC during a migrant phase.

The **no** form of this command reverts to the default.

Default initial-valid-lifetime min 5

Parameters *hours* — Specifies the number of initial preferred lifetime hours.

Values 1 to 1

minutes — Specifies the number of initial preferred lifetime minutes.

Values 5 to 59

seconds — Specifies the number of initial preferred lifetime seconds.

Values 1 to 59

Combined values: min 5 – hrs 1

initial-lease-time

Syntax **initial-lease-time** [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no initial-lease-time

Context config>service>vprn>sub-if>grp-if>wlan-gw>dhcp
config>service>ies>sub-if>grp-if>wlan-gw>dhcp
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dhcp
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dhcp

Description This command configures the lease time for a user which is migrant (unauthenticated).

Default initial-lease-time min 10

Parameters *hours* — Specifies the number of initial lease time hours.

Values 1 to 1

minutes — Specifies the number of initial lease time minutes.

Values 5 to 59

seconds — Specifies the number of initial lease time.

Values 1 to 59

l2-aware-ip-address

Syntax **l2-aware-ip-address** *ip-address*
l2-aware-ip-address from-pool
no l2-aware-ip-address

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dhcp
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dhcp

Description This command configures the L2-Aware NAT inside IP address to be assigned via DHCP on WLAN-GW ISA.

If the **from-pool** parameter is specified instead of an IPv4 address, a unique address is allocated to each UE. The pool used is managed by the dhcpv4-nat pool manager, configured under the same subscriber interface. This option is only available when **auth-on-dhcp** is also configured.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the L2-Aware NAT inside IP address.

from-pool — Specifies that the L2-Aware IP address is allocated from a pool.

primary-dns

Syntax **primary-dns** *ip-address*
no primary-dns

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dhcp
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dhcp

Description This command configures the primary DNS address to be returned via DHCP on WLAN-GW ISA.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the primary DNS address.

secondary-dns

Syntax **secondary-dns** *ip-address*
no secondary-dns

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dhcp
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dhcp

Description This command configures the secondary DNS address to be returned via DHCP on WLAN-GW ISA.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the secondary DNS address.

primary-nbns

Syntax **primary-nbns** *ip-address*

no primary-nbns

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dhcp
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dhcp

Description This command configures the primary NBNS address to be returned via DHCP on WLAN-GW ISA.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the primary NBNS address.

secondary-nbns

Syntax **secondary-nbns** *ip-address*
no secondary-nbns

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>dhcp
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>dhcp

Description This command configures the secondary NBNS address to be returned via DHCP on WLAN-GW ISA.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the secondary NBNS address.

idle-timeout

Syntax **idle-timeout action** *idle-timeout-action*
no idle-timeout

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range
config>service>ies >sub-if>grp-if>wlan-gw>ranges>range

Description This command specifies idle-timeout behavior for DSM UEs and UEs undergoing (ISA-based) portal authentication. This knob only specifies the desired action, idle-timeout is activated by RADIUS on a per-UE basis.

The **no** form of this command resets the idle-timeout to its default

Default idle-timeout action remove

Parameters *action* — Specifies which action to perform when the idle-timeout timer goes off.

Values remove, shcv

http-redirect-policy

Syntax **http-redirect-policy** *policy-name*
no http-redirect-policy

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range

Description This command specifies http redirect policy on ISA to redirect http traffic to the URL specified in the policy.

The **no** form of this command reverts to the default.

Parameters *policy-name* — Specifies the name of the http redirect policy under subscriber-management context.

l2-service

Syntax **l2-service** *service-id*
no l2-service

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range
config>service>ies >sub-if>grp-if>wlan-gw>ranges>range

Description This command specifies the VPLS service used for L2 wholesale. When such a service is configured no other configuration is allowed under the vlan-range.

The **no** form of this command removes the L2 wholesale service, this is only allowed if the l2-service node is shutdown.

Parameters *service-id* — Specifies the VPLS service ID to use for Layer 2 wholesale.

nat-policy

Syntax **nat-policy** *policy-name*
no nat-policy

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range

Description This command specifies the NAT policy for WLAN-GW ISA subscribers.

The **no** form of this command reverts to the default.

brg

Syntax **brg**

Context config>service>ies>sub-if>grp-if>wlan-gw>ranges>range
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range
config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command enables the context to configure BRG parameters. In the **config>service>ies>sub-if>grp-if** and **config>service>vprn>sub-if>grp-if** contexts, these commands are only available in the **vlan-tag-ranges** context.

authenticated-brg-only

Syntax [**no**] **authenticated-brg-only**

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>brg
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>brg
config>service>vprn>sub-if>grp-if>brg
config>service>ies>sub-if>grp-if>brg

Description This command indicates that only BRGs that are pre-authenticated using the RADIUS proxy are allowed in this context.

The **no** form of this command removes the restriction

default-brg-profile

Syntax **default-brg-profile** *profile-name*
no default-brg-profile

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>brg
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>brg
config>service>vprn>sub-if>grp-if>brg
config>service>ies>sub-if>grp-if>brg

Description This command indicates that the default BRG profile must be used for new BRGs. This profile can be overridden by RADIUS.

The **no** form of this command removes the profile name from the configuration.

Parameters *profile-name* — Specifies the name of the brg-profile to be applied.

data-triggered-ue-creation

Syntax [**no**] **data-triggered-ue-creation**

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range

Description This command enables or disables data-triggered subscriber creation for WIFI subscribers. Data triggered UE creation is currently only supported for UDP and TCP packets.

The **no** form of this command reverts to the default.

track-mobility

Syntax **track-mobility**

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range

Description This command enables the context to configure RADIUS-proxy cache information required for subscribers that are created via data-triggered authentication. The RADIUS proxy cache enables efficient handling of UE mobility.

mac-format

Syntax **mac-format** *mac-format*
no mac-format

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>track-mobility
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>track-mobility

Description This command configures how the MAC address is represented by the RADIUS proxy server.

Default no mac-format "aa:"

Parameters *mac-format* — Specifies how the MAC address is represented by the RADIUS proxy server.

Values

mac-format	like ab: for 00:0c:f1:99:85:b8
	or XY- for 00-0C-F1-99-85-B8
	or mmmm. for 0002.03aa.abff
	or xx for 000cf19985b8

radius-proxy-cache

Syntax **radius-proxy-cache router** *router-instance* **server** *server-name*
no radius-proxy-cache

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>track-mobility
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>track-mobility

Description This command specifies the RADIUS-proxy server to allow subscribers created via data-triggered authentication to create an entry. This RADIUS proxy cache entry allows efficient handling of UE mobility.

Parameters *router-instance* — Specifies the router instance.

Values

router-name	Base
service-id	1 to 2147483647

server-name — Specifies the server name up to 32 characters.

sap-template

Syntax **sap-template** *sap-template*
no sap-template

Context config>service>vpls>wlan-gw

Description This command specifies the VPLS SAP template that is applied on the internal SAPs created for communication between the VPLS and the ISAs.

The **no** form of this command removes the SAP template.

Parameters *sap-template* — Specifies the existing SAP template to apply. The template is created in the **config>service>template** context.

11.25.2.8.1 Show Commands

Note: The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

call-trace

Syntax **call-trace**

Context show

Description This command enables the context to display information related to the call-trace module.

wlan-gw

Syntax **wlan-gw**

Context show>call-trace

Description The command enables the context to display information related to the wlan-gw call-trace functionality.

ue

Syntax **ue** [*ieee-address*] [**detail**]

Context show>call-trace

Description This command gives an overview of either all traces or a specific trace on the WLAN-GW.

Parameters *ieee-address* — Displays information about the MAC address of this UE

detail — Displays detailed information about the job

Output The following output is an example of traces of the UE being monitored.

Sample Output

```
Node# show call-trace hosts
===============================================================================
Call-trace hosts
===============================================================================
 MAC address Mask-name     Status      Msgs
-------------------------------------------------------------------------------
 00:0a:95:9d:68:16 N/A               running   16
-------------------------------------------------------------------------------
Number of call-trace debug jobs: 1
===============================================================================
Node# show call-trace hosts detail
===============================================================================
Call-trace  detail
===============================================================================
MAC address                          : 00:0a:95:95:34:0a                Status :
running
                  Capture format     : pcap
Nr. of captured msgs   : 4 Time limit            : 86400s
Size of captured msgs : 2620B Data limit             : 10MB
Started : NOV 12 2013, 15:28:17 UTC
Live output : N/A
-------------------------------------------------------------------------------
```

acct-on-off-group

Syntax **acct-on-off-group** [*group-name*]

Context show>aaa

Description This command displays Acct-On-Off group information and the associated RADIUS server policies.

Parameters *group-name* — Displays information pertaining to the specified acct-on-off group

Output The following output is an example of AAA Acct-On-Off group information.

Sample Output

```
# show aaa acct-on-off-group "group-1"
===============================================================================
Acct-On-Off-Group Information
===============================================================================
acct on off group name               : group-1
  - controlling Radius-Server-policy :
        aaa-server-policy-3
  - monitored by Radius-Server-policy :
        aaa-server-policy-4
-------------------------------------------------------------------------------
Nbr of Acct-on-off-groups displayed : 1
-------------------------------------------------------------------------------
===============================================================================
```

Table 164 **WiFi Acct-On-Off Field Descriptions**

Label	Description
acct on off group name	The name of a RADIUS server policy Accounting-On-Off-Group
controlling Radius-Server-policy	The controlling RADIUS server policy name
monitored by Radius-Server-policy	The policy monitored a RADIUS server policy.
Nbr of Acct-on-off-groups displayed	The RADIUS policy that controls the Acct-On-Off group

radius-proxy-server

Syntax **radius-proxy-server** *server-name*
radius-proxy-server *server-name* **cache**

radius-proxy-server *server-name* **cache hex-key** *hex-string*
radius-proxy-server *server-name* **cache string-key** *string*
radius-proxy-server *server-name* **cache summary**
radius-proxy-server *server-name* **statistics**
radius-proxy-server

Context show>router

Description This command displays summary of RADIUS-proxy cache or specific entries.

Parameters *server-name* — Displays information about the specified server name

cache — Displays messages used to generate the key for the cache of this RADIUS proxy server

hex-key *hex-string* — Displays information about the specified hex string

Values 0x0 to 0xFFFFFFFF (maximum of 64 hex nibbles)]

string — Displays information about the specified string

summary — Displays a summary of the cache of the RADIUS proxy servers

statistics — Displays statistics about the RADIUS proxy servers of this system

Output The following is an example of RADIUS proxy server information.

Sample Output

```
system# show router 10 radius-proxy-server "myProxyServer1"
===============================================================================
RADIUS Proxy server "myProxyServer1"
===============================================================================
Description                 : myDesc
Purpose                     : authentication
Administrative state        : in-service
Default acct server policy  : myRadiusServerPolicy1
Default auth server policy  : myRadiusServerPolicy2
Send accounting response    : true
Last management change      : 02/17/2012 14:54:28
-------------------------------------------------------------------------------
Cache settings
-------------------------------------------------------------------------------
Administrative state        : enabled
Key packet type             : access-accept
Key attribute type          : 12
Key vendor ID               : (Not Specified)
Timeout (s)                 : 60
Track accounting            : stop interim-update accounting-on accounting-off
Load balance key            : source-ip-udp
===============================================================================
Interfaces
-------------------------------------------------------------------------------
myInterface1
myInterface2
myInterface3
-------------------------------------------------------------------------------
```

```
No. of Interface(s): 3
===============================================================================
Usernames/RADIUS server policies
===============================================================================
Id Username-match                    RADIUS-server-policy            Purpose
-------------------------------------------------------------------------------
1.  aaa                               myRadiusServerPolicy2 auth
===============================================================================
system#
```

Table 165 RADIUS Proxy Server Field Descriptions

Label	Description
Description	The description of this RADIUS proxy server
Purpose	The purpose of the RADIUS server, either accounting or authentication
Administrative state	The administrative state of this RADIUS server
Default acct server policy	The name of the default RADIUS server policy associated with this RADIUS proxy server for accounting purposes
Default auth server policy	The name of the default RADIUS server policy associated with this RADIUS proxy server for authentication purposes
Send accounting response	Specifies if this RADIUS Proxy server itself responds with an Accounting-Response message to each received Accounting-Request instead of proxying them to a configured RADIUS server
Last management change	The sysUpTime at the time of the most recent management-initiated change
Key packet type	The packet type of the RADIUS messages to use to generate the key for the cache of this RADIUS proxy server, access-request, access-accept, access-reject, access-challenge
Key attribute type	The RADIUS attribute type to cache for this RADIUS proxy server. Refer to RFC 2865, *Remote Authentication Dial In User Service (RADIUS*), *Section 5 Attributes*.
Key vendor ID	The RADIUS Vendor-Id. Refer to RFC 2865, *Remote Authentication Dial In User Service (RADIUS), Section 5.25 Vendor-Specific*
Timeout (s)	Displays, in seconds, the timeout after which an entry in the cache will expire
Track accounting	The RADIUS accounting packets that have impact on the cache of this RADIUS proxy server

Table 165 **RADIUS Proxy Server Field Descriptions (Continued)**

Label	Description
Load balance key	The key for load-balancing RADIUS messages between RADIUS servers
Id	The specifies the RADIUS Vendor-Id
Username	The user name
RADIUS-server-policy	The RADIUS server name
Purpose	The purpose of the RADIUS server, either accounting or authentication

wlan-gw

Syntax **wlan-gw**

Context show>router

Description This command displays Wireless LAN Gateway information.

isa-subnets

Syntax **isa-subnets** [**detail**]
isa-subnets [**detail**] **interface** *interface-name*
isa-subnets prefix *ipv6-address/prefix-length*

Context show>router>wlan-gw

Description This command outputs all the prefixes in use by the WLAN GW pool manager.

Parameters **detail** — Displays detailed information for each prefix

interface-name — Displays only the prefixes associated with this subscriber interface

ipv6-address/prefix-length — Displays details of a specific IPv6 address and prefix

Output The following is an example of WLAN-GW ISA subnet information.

Sample Output

```
system# show router wlan-gw isa-subnets
===============================================================================
ISA Subnets
===============================================================================
Prefix                                                  MDA     Family  Usage
-------------------------------------------------------------------------------
```

```
2001:db8:0:1/48                                          3/1     dhcpv6  0%
2001:db8:1::/48                                        3/2     dhcpv6  0%
2001:db8:2::/48                                        4/1     dhcpv6  0%
2001:db8:3::/48                                        4/2     dhcpv6  0%
2001:db8:4::/48                                        5/1     dhcpv6  0%
2001:db8:5::/48                                        5/2     dhcpv6  0%
2001:db8:6::/48                                        3/1     slaac   0%
2001:db8:7::/48                                        3/2     slaac   0%
2001:db8:8::/48                                        4/1     slaac   0%
2001:db8:9::/48                                        4/2     slaac   0%
2001:db8:a::/48                                        5/1     slaac   0%
2001:db8:b::/48                                        5/2     slaac   0%
-------------------------------------------------------------------------------
No. of ISA subnets: 12
===============================================================================

*A:Dut-C# show router wlan-gw isa-subnets prefix 2001:db8::/48
===============================================================================
ISA Subnet Prefix           : 2001:db8::/48
-------------------------------------------------------------------------------
Group Id                    : 1
Member Id                   : 1
MDA                         : 3/1
Family                      : dhcpv6
Subscriber Interface        : wlangw-sub-itf
Pool Is Old                 : No
Usage Level                 : 0%
Remaining Lease Time        : 0d 23:50:54
DHCPv6 Options              : (length=512)
                            : 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```

```
                                : 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
                                : 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
===============================================================================
```

mgw-address-cache

Syntax **mgw-address-cache** [**arec**] [**snaptr**] [**srv**]
mgw-address-cache apn *apn-domain-string*

Context show>router>wlan-gw

Description This command displays the mobile gateway's DNS lookup address cache.

Parameters **arec** — Displays A-records

snaptr — Displays Straightforward-NAPTR information

srv — Displays SRV records

apn-domain-string — Specifies the Access Point Name (APN) of this DNS cache entry

Output The following output is an example of WLAN-GW DNS lookup address cache information.

Sample Output

```
*A:Dut-C# show router 300 wlan-gw mgw-address-cache
===============================================================================
Mobile Gateway SNAPTR cache
===============================================================================
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.
                              3gppnetwork.org
Order                       : 10
Index                       : 1
-------------------------------------------------------------------------------
Preference                  : 10
Service                     : x-3gpp-pgw:x-gn-gtp
Next lookup                 : dns-srv
Replacement                 : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.
                              3gppnetwork.org
Time left (s)               : 3582
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.
                              3gppnetwork.org
Order                       : 20
Index                       : 2
-------------------------------------------------------------------------------
Preference                  : 20
Service                     : x-3gpp-pgw:x-s2a-gtp
Next lookup                 : dns-srv
Replacement                 : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.
                              3gppnetwork.org
Time left (s)               : 3582
-------------------------------------------------------------------------------
```

```
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.
                              3gppnetwork.org
Order                       : 30
Index                       : 3
--------------------------------------------------------------------------------
Preference                  : 30
Service                     : x-3gpp-pgw:x-s2b-gtp
Next lookup                 : dns-srv
Replacement                 : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.
                              3gppnetwork.org
Time left (s)               : 3581
--------------------------------------------------------------------------------
No. of SNAPTR cache entries: 3
================================================================================
================================================================================
Mobile Gateway SRV cache
================================================================================
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.
                              3gppnetwork.org
Priority                    : 10
Index                       : 1
--------------------------------------------------------------------------------
Weight                      : 10
Port                        : 2123
Target                      : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a1.
                              3gppnetwork.org
Time left (s)               : 3581
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.
                              3gppnetwork.org
Priority                    : 20
Index                       : 2
--------------------------------------------------------------------------------
Weight                      : 20
Port                        : 2123
Target                      : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a2.
                              3gppnetwork.org
Time left (s)               : 3581
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.
                              3gppnetwork.org
Priority                    : 10
Index                       : 1
--------------------------------------------------------------------------------
Weight                      : 10
Port                        : 2123
Target                      : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a1.
                              3gppnetwork.org
Time left (s)               : 3581
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.
                              3gppnetwork.org
Priority                    : 20
Index                       : 2
--------------------------------------------------------------------------------
Weight                      : 20
Port                        : 2123
Target                      : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a2.
```

```
                                3gppnetwork.org
Time left (s)               : 3581
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.
                              3gppnetwork.org
Priority                    : 10
Index                       : 1
-------------------------------------------------------------------------------
Weight                      : 10
Port                        : 2123
Target                      : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a1.
                              3gppnetwork.org
Time left (s)               : 3581
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.
                              3gppnetwork.org
Priority                    : 20
Index                       : 2
-------------------------------------------------------------------------------
Weight                      : 20
Port                        : 2123
Target                      : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a2.
                              3gppnetwork.org
Time left (s)               : 3581
-------------------------------------------------------------------------------
No. of SRV cache entries: 6
===============================================================================
===============================================================================
Mobile Gateway address cache
===============================================================================
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a1.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10.0.0.23
Time left (s)               : 3581
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a1.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10.0.0.29
Time left (s)               : 3581
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a1.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      :  10.0.0.35
Time left (s)               : 3581
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a2.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      :  10.0.0.24
Time left (s)               : 3581
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a2.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.30
```

```
Time left (s)               : 3581
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a2.
                              3gppnetwork.org
--------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.36
Time left (s)               : 3581
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a1.
                              3gppnetwork.org
--------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.25
Time left (s)               : 3581
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a1.
                              3gppnetwork.org
--------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.31
Time left (s)               : 3581
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a1.
                              3gppnetwork.org
--------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.37
Time left (s)               : 3581
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a2.
                              3gppnetwork.org
--------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.26
Time left (s)               : 3580
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a2.
                              3gppnetwork.org
--------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.32
Time left (s)               : 3580
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a2.
                              3gppnetwork.org
--------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.38
Time left (s)               : 3580
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a1.
                              3gppnetwork.org
--------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.27
Time left (s)               : 3580
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a1.
                              3gppnetwork.org
--------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.33
Time left (s)               : 3580
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a1.
                              3gppnetwork.org
--------------------------------------------------------------------------------
```

```
Mobile Gateway address      : 10,0.0.39
Time left (s)               : 3580
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a2.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.28
Time left (s)               : 3580
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a2.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.34
Time left (s)               : 3580
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a2.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.40
Time left (s)               : 3580
-------------------------------------------------------------------------------
No. of cache entries: 18


*A:Dut-C# show router 300 wlan-gw mgw-address-cache arec
===============================================================================
Mobile Gateway address cache
===============================================================================
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a1.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.23
Time left (s)               : 3573
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a1.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.29
Time left (s)               : 3573
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a1.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.35
Time left (s)               : 3573
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a2.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.24
Time left (s)               : 3573
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a2.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.30
Time left (s)               : 3573
-------------------------------------------------------------------------------
```

```
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a2.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.36
Time left (s)               : 3573
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a1.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.25
Time left (s)               : 3573
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a1.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.31
Time left (s)               : 3573
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a1.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.37
Time left (s)               : 3573
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a2.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.26
Time left (s)               : 3573
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a2.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.32
Time left (s)               : 3573
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a2.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.38
Time left (s)               : 3572
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a1.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.27
Time left (s)               : 3572
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a1.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.33
Time left (s)               : 3572
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a1.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.39
Time left (s)               : 3572
```

```
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a2.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.28
Time left (s)               : 3572
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a2.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.34
Time left (s)               : 3572
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a2.
                              3gppnetwork.org
-------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.40
Time left (s)               : 3572
-------------------------------------------------------------------------------
No. of cache entries: 18
===============================================================================


*A:Dut-C# show router 300 wlan-gw mgw-address-cache srv
===============================================================================
Mobile Gateway SRV cache
===============================================================================
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.
                              3gppnetwork.org
Priority                    : 10
Index                       : 1
-------------------------------------------------------------------------------
Weight                      : 10
Port                        : 2123
Target                      : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a1.
                              3gppnetwork.org
Time left (s)               : 3567

-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.
                              3gppnetwork.org
Priority                    : 20
Index                       : 2
-------------------------------------------------------------------------------
Weight                      : 20
Port                        : 2123
Target                      : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a2.
                              3gppnetwork.org
Time left (s)               : 3567

-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.
                              3gppnetwork.org
Priority                    : 10
Index                       : 1
-------------------------------------------------------------------------------
Weight                      : 10
Port                        : 2123
```

```
Target                      : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a1.
                              3gppnetwork.org
Time left (s)               : 3566

-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.
                              3gppnetwork.org
Priority                    : 20
Index                       : 2
-------------------------------------------------------------------------------
Weight                      : 20
Port                        : 2123
Target                      : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a2.
                              3gppnetwork.org
Time left (s)               : 3566
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.
                              3gppnetwork.org
Priority                    : 10
Index                       : 1
-------------------------------------------------------------------------------
Weight                      : 10
Port                        : 2123
Target                      : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a1.
                              3gppnetwork.org
Time left (s)               : 3566
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.
                              3gppnetwork.org
Priority                    : 20
Index                       : 2
-------------------------------------------------------------------------------
Weight                      : 20
Port                        : 2123
Target                      : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a2.
                              3gppnetwork.org
Time left (s)               : 3566
-------------------------------------------------------------------------------
No. of SRV cache entries: 6
===============================================================================


*A:Dut-C# show router 300 wlan-gw mgw-address-cache snaptr
===============================================================================
Mobile Gateway SNAPTR cache
===============================================================================
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.
                              3gppnetwork.org
Order                       : 10
Index                       : 1
-------------------------------------------------------------------------------
Preference                  : 10
Service                     : x-3gpp-pgw:x-gn-gtp
Next lookup                 : dns-srv
Replacement                 : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.
                              3gppnetwork.org
Time left (s)               : 3555
-------------------------------------------------------------------------------
```

```
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.
                              3gppnetwork.org
Order                       : 20
Index                       : 2
-------------------------------------------------------------------------------
Preference                  : 20
Service                     : x-3gpp-pgw:x-s2a-gtp
Next lookup                 : dns-srv
Replacement                 : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.
                              3gppnetwork.org
Time left (s)               : 3555
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.
                              3gppnetwork.org
Order                       : 30
Index                       : 3
-------------------------------------------------------------------------------
Preference                  : 30
Service                     : x-3gpp-pgw:x-s2b-gtp
Next lookup                 : dns-srv
Replacement                 : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.
                              3gppnetwork.org
Time left (s)               : 3554

-------------------------------------------------------------------------------
No. of SNAPTR cache entries: 3


*A:Dut-C# show router 300 wlan-gw mgw-address-
cache apn full.dotted.apn.apn.epc.mnc010.mcc206.3gppnetwork.org
===============================================================================
Mobile Gateway APN Cache
===============================================================================
-------------------------------------------------------------------------------
APN > NAPTR
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.
                              3gppnetwork.org
Order                       : 10
Index                       : 1
-------------------------------------------------------------------------------
Preference                  : 10
Service                     : x-3gpp-pgw:x-gn-gtp
Next lookup                 : dns-srv
Replacement                 : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.
                              3gppnetwork.org
Time left (s)               : 3531
-------------------------------------------------------------------------------
APN > NAPTR > SRV
-------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.
                              3gppnetwork.org
Priority                    : 10
Index                       : 1
-------------------------------------------------------------------------------
Weight                      : 10
Port                        : 2123
Target                      : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a1.
                              3gppnetwork.org
```

```
Time left (s)               : 3531
--------------------------------------------------------------------------------
APN > NAPTR > SRV > A
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a1.
                              3gppnetwork.org
--------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.23
Time left (s)               : 3530
--------------------------------------------------------------------------------
APN > NAPTR > SRV > A
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a1.
                              3gppnetwork.org
--------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.29
Time left (s)               : 3530
--------------------------------------------------------------------------------
APN > NAPTR > SRV > A
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a1.
                              3gppnetwork.org
--------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.35
Time left (s)               : 3530
--------------------------------------------------------------------------------
APN > NAPTR > SRV
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.
                              3gppnetwork.org
Priority                    : 20
Index                       : 2
--------------------------------------------------------------------------------
Weight                      : 20
Port                        : 2123
Target                      : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a2.
                              3gppnetwork.org
Time left (s)               : 3530
--------------------------------------------------------------------------------
APN > NAPTR > SRV > A
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a2.
                              3gppnetwork.org
--------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.24
Time left (s)               : 3530
--------------------------------------------------------------------------------
APN > NAPTR > SRV > A
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a2.
                              3gppnetwork.org
--------------------------------------------------------------------------------
Mobile Gateway address      : 10,0.0.30
Time left (s)               : 3530
--------------------------------------------------------------------------------
APN > NAPTR > SRV > A
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv1.a2.
                              3gppnetwork.org
```

```
--------------------------------------------------------------------------------
Mobile Gateway address      : 10.0.0.36
Time left (s)               : 3530
--------------------------------------------------------------------------------
APN > NAPTR
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.
                              3gppnetwork.org
Order                       : 20
Index                       : 2
--------------------------------------------------------------------------------
Preference                  : 20
Service                     : x-3gpp-pgw:x-s2a-gtp
Next lookup                 : dns-srv
Replacement                 : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.
                              3gppnetwork.org
Time left (s)               : 3530
--------------------------------------------------------------------------------
APN > NAPTR > SRV
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.
                              3gppnetwork.org
Priority                    : 10
Index                       : 1
--------------------------------------------------------------------------------
Weight                      : 10
Port                        : 2123
Target                      : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a1.
                              3gppnetwork.org
Time left (s)               : 3529
--------------------------------------------------------------------------------
APN > NAPTR > SRV > A
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a1.
                              3gppnetwork.org
--------------------------------------------------------------------------------
Mobile Gateway address      : 10.0.0.25
Time left (s)               : 3529
--------------------------------------------------------------------------------
APN > NAPTR > SRV > A
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a1.
                              3gppnetwork.org
--------------------------------------------------------------------------------
Mobile Gateway address      : 10.0.0.31
Time left (s)               : 3529
--------------------------------------------------------------------------------
APN > NAPTR > SRV > A
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a1.
                              3gppnetwork.org
--------------------------------------------------------------------------------
Mobile Gateway address      : 10.0.0.37
Time left (s)               : 3529
--------------------------------------------------------------------------------
APN > NAPTR > SRV
--------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.
                              3gppnetwork.org
```

```
Priority                    : 20
Index                       : 2
------------------------------------------------------------------------------
Weight                      : 20
Port                        : 2123
Target                      : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a2.
                              3gppnetwork.org
Time left (s)               : 3529
------------------------------------------------------------------------------
APN > NAPTR > SRV > A
------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a2.
                              3gppnetwork.org
------------------------------------------------------------------------------
Mobile Gateway address      : 10.0.0.26
Time left (s)               : 3529
------------------------------------------------------------------------------
APN > NAPTR > SRV > A
------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a2.
                              3gppnetwork.org
------------------------------------------------------------------------------
Mobile Gateway address      : 10.0.0.32
Time left (s)               : 3529

------------------------------------------------------------------------------
APN > NAPTR > SRV > A
------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv2.a2.
                              3gppnetwork.org
------------------------------------------------------------------------------
Mobile Gateway address      : 10.0.0.38
Time left (s)               : 3529
------------------------------------------------------------------------------
APN > NAPTR
------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.
                              3gppnetwork.org
Order                       : 30
Index                       : 3
------------------------------------------------------------------------------
Preference                  : 30
Service                     : x-3gpp-pgw:x-s2b-gtp
Next lookup                 : dns-srv
Replacement                 : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.
                              3gppnetwork.org
Time left (s)               : 3529
------------------------------------------------------------------------------
APN > NAPTR > SRV
------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.
                              3gppnetwork.org
Priority                    : 10
Index                       : 1
------------------------------------------------------------------------------
Weight                      : 10
Port                        : 2123
Target                      : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a1.
                              3gppnetwork.org
```

```
Time left (s)               : 3529
------------------------------------------------------------------------------
APN > NAPTR > SRV > A
------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a1.
                              3gppnetwork.org
------------------------------------------------------------------------------
Mobile Gateway address      : 10.0.0.27
Time left (s)               : 3529
------------------------------------------------------------------------------
APN > NAPTR > SRV > A
------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a1.
                              3gppnetwork.org
------------------------------------------------------------------------------
Mobile Gateway address      : 10.0.0.33
Time left (s)               : 3529
------------------------------------------------------------------------------
APN > NAPTR > SRV > A
------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a1.
                              3gppnetwork.org
------------------------------------------------------------------------------
Mobile Gateway address      : 10.0.0.39
Time left (s)               : 3529
------------------------------------------------------------------------------
APN > NAPTR > SRV
------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.
                              3gppnetwork.org
Priority                    : 20
Index                       : 2
------------------------------------------------------------------------------
Weight                      : 20
Port                        : 2123
Target                      : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a2.
                              3gppnetwork.org
Time left (s)               : 3529
------------------------------------------------------------------------------
APN > NAPTR > SRV > A
------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a2.
                              3gppnetwork.org
------------------------------------------------------------------------------
Mobile Gateway address      : 10.0.0.28
Time left (s)               : 3529
------------------------------------------------------------------------------
APN > NAPTR > SRV > A
------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a2.
                              3gppnetwork.org
------------------------------------------------------------------------------
Mobile Gateway address      : 10.0.0.34
Time left (s)               : 3528
------------------------------------------------------------------------------
APN > NAPTR > SRV > A
------------------------------------------------------------------------------
APN                         : full.dotted.apn.apn.epc.mnc010.mcc206.srv3.a2.
                              3gppnetwork.org
```

```
-------------------------------------------------------------------------------
Mobile Gateway address      : 10.0.0.40
Time left (s)               : 3528
-------------------------------------------------------------------------------
No. of cache entries: 18
```

tunnel-qos

Syntax **tunnel-qos** [**detail**]
tunnel-qos remote-ip *ip-address* [**local-ip** *ip-address*] [**detail**]

Context show>router>wlan-gw

Description This command displays tunnel-QoS resource information.

Parameters *ip-address* — Specifies the IP address of the Mobile Gateway that is the source IP address in the tunnel header of received packets

Values

ip-address:	ipv4-address - a.b.c.d
ipv6-address :	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D

ip-address — Specifies the IP address of this system that is the destination IP address in the tunnel header of received packets

Values

ip-address:	ipv4-address - a.b.c.d
ipv6-address :	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D

detail — Displays detailed information.

Output The following output is an example of soft GRE tunnel QoS information.

Sample Output

```
*A:Dut-C# show router 50 wlan-gw soft-gre-tunnel-qos
===============================================================================
Soft GRE tunnel QoS
===============================================================================
Remote IP address           : 239.0.0.2
Local IP address            : 10.1.1.1
Operational state           : active
```

```
Number of UE                : 1
Remaining hold time (s)     : N/A


*A:Dut-C# show router 50 wlan-gw soft-gre-tunnel-qos detail
===============================================================================
Soft GRE tunnel QoS
===============================================================================
Remote IP address           : 239.0.0.2
Local IP address            : 10.1.1.1
Operational state           : active
Number of UE                : 1
Remaining hold time (s)     : N/A
Service Access Points(SAP)
===============================================================================
Service Id         : 2147483650
SAP                : 5/1/lo-gre:1             Encap             : q-tag
Description        : Internal SAP
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 03/24/2014 15:03:48
Last Mgmt Change   : 03/24/2014 15:14:00

-------------------------------------------------------------------------------
Encap Group Specifics
-------------------------------------------------------------------------------
Encap Group Name   : _tmnx_SHAPER_GR000       Group Type        : ISID
Qos-per-member     : TRUE
Members            :
1

-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
E. qos-policy      : 1                        Q Frame-Based Acct: Disabled
E. Sched Policy    :                          E. Agg-limit      : -1
                                              Limit Unused BW   : Disabled
-------------------------------------------------------------------------------
Encap Group Member 1 Base Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

Forwarding Engine Stats
                        Packets                 Octets

For. InProf           : 0                       0
For. OutProf          : 0                       0
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0

-------------------------------------------------------------------------------
Encap Group Member 1 Queue Statistics
-------------------------------------------------------------------------------

                        Packets                 Octets

Egress Queue 1
For. InProf           : 0                       0
```

```
For. OutProf          : 0                       0
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0


*A:Dut-C# show router 50 wlan-gw soft-gre-tunnel-qos remote-ip 239.0.0.2
===============================================================================
Soft GRE tunnel QoS
===============================================================================
Remote IP address           : 239.0.0.2
Local IP address            : 10.1.1.1
Operational state           : active
Number of UE                : 1
Remaining hold time (s)     : N/A


*A:Dut-C# show router 50 wlan-gw soft-gre-tunnel-qos  remote-ip 239.0.0.2  local-
ip 10.1.1.1
===============================================================================
Soft GRE tunnel QoS
===============================================================================
Remote IP address           : 239.0.0.2
Local IP address            : 10.1.1.1
Operational state           : active
Number of UE                : 1
Remaining hold time (s)     : N/A


*A:Dut-C# show router 50 wlan-gw soft-gre-tunnel-qos  remote-ip 239.0.0.2  local-
ip 10.1.1.1 detail
===============================================================================
Soft GRE tunnel QoS
===============================================================================
Remote IP address           : 239.0.0.2
Local IP address            : 10.1.1.1
Operational state           : active
Number of UE                : 1
Remaining hold time (s)     : N/A
Service Access Points(SAP)
===============================================================================
Service Id         : 2147483650
SAP                : 5/1/lo-gre:1             Encap             : q-tag
Description        : Internal SAP
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 03/24/2014 15:03:48
Last Mgmt Change   : 03/24/2014 15:14:00
-------------------------------------------------------------------------------
Encap Group Specifics
-------------------------------------------------------------------------------
Encap Group Name   : _tmnx_SHAPER_GR000       Group Type        : ISID
Qos-per-member     : TRUE
Members            :
1
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
E. qos-policy      : 1                        Q Frame-Based Acct: Disabled
```

```
E. Sched Policy    :                          E. Agg-limit      : -1
                                              Limit Unused BW   : Disabled
-------------------------------------------------------------------------------
Encap Group Member 1 Base Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

Forwarding Engine Stats
                        Packets                 Octets

For. InProf           : 0                       0
For. OutProf          : 0                       0
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
-------------------------------------------------------------------------------
Encap Group Member 1 Queue Statistics
-------------------------------------------------------------------------------
                        Packets                 Octets

Egress Queue 1
For. InProf           : 0                       0
For. OutProf          : 0                       0
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
```

tunnels

Syntax **tunnels local-ip** *ip-address* **remote-ip** *ip-address* **ue**
tunnels [**local-ip** *ip-address*] [**remote-ip** *ip-address*] [**isa-group** *wlan-gw-group-id*] [**member** [1 to 255]] [**summary**] [**detail**]

Context show>router>wlan-gw

Description This command displays tunnel-QoS resource information.

Parameters *ip-address* — Specifies the remote IP address of the Mobile Gateway that is the source IP address in the tunnel header of received packets

Values

ip-address:	ipv4-address - a.b.c.d
ipv6-address :	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D

ip-address — Specifies the local IP address of this system that is the destination IP address in the tunnel header of received packets

Values

ip-address:	ipv4-address - a.b.c.d

ipv6-address : x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x - [0 to FFFF]H
d - [0 to 255]D

ue — Displays information for the specified user equipment

wlan-gw-group-id — Specifies the identifier of the WLAN Gateway ISA group that terminates GRE for this group interface

Values 1 to 4

member — Specifies the identifier of this WLAN Gateway ISA group member

Values 1 to 255

summary — Displays a summary of the specified parameters

detail — Displays detailed information

Output The following output is an example of WLAN-GW soft GRE tunnel information.

Sample Output

```
*A:Dut-C# show router 50 wlan-gw soft-gre-tunnels
===============================================================================
Soft GRE tunnels
===============================================================================
Remote IP address           : 239.0.0.2
Local IP address            : 10.1.1.1
ISA group ID                : 1
ISA group member ID         : 5
Time established            : 2014/03/24 15:38:52
Number of UE                : 1
Access Point MAC            : 00:00:00:00:00:01
AP MAC learn failed         : false

Tunnel QoS
----------
Operational state           : active
Number of UE                : 1
Remaining hold time (s)     : N/A

-------------------------------------------------------------------------------
No. of tunnels: 1

*A:Dut-C# show router 50 wlan-gw soft-gre-tunnels local-ip 10.1.1.1
===============================================================================
Soft GRE tunnels
===============================================================================
Remote IP address           : 239.0.0.2
Local IP address            : 10.1.1.1
ISA group ID                : 1
ISA group member ID         : 5
Time established            : 2014/03/24 15:38:52
Number of UE                : 1
```

```
Access Point MAC            : 00:00:00:00:00:01
AP MAC learn failed         : false

Tunnel QoS
----------
Operational state           : active
Number of UE                : 1
Remaining hold time (s)     : N/A
-------------------------------------------------------------------------------
No. of tunnels: 1

*A:Dut-C# show router 50 wlan-gw soft-gre-tunnels local-ip 10.1.1.1 remote-
ip 239.0.0.2
===============================================================================
Soft GRE tunnels
===============================================================================
Remote IP address           : 239.0.0.2
Local IP address            : 10.1.1.1
ISA group ID                : 1
ISA group member ID         : 5
Time established            : 2014/03/24 15:38:52
Number of UE                : 1
Access Point MAC            : 00:00:00:00:00:01
AP MAC learn failed         : false

Tunnel QoS
----------
Operational state           : active
Number of UE                : 1
Remaining hold time (s)     : N/A
-------------------------------------------------------------------------------
No. of tunnels: 1

*A:Dut-C# show router 50 wlan-gw soft-gre-tunnels local-ip 10.1.1.1 remote-
ip 239.0.0.2 ue
===============================================================================
Tunnel User Equipments
===============================================================================
MAC address                 : 00:02:00:00:00:01
-------------------------------------------------------------------------------
VLAN Q-tag                  : 1
MPLS label                  : (Not Specified)
Tunnel router               : 50
Tunnel remote IP address    : 239.0.0.2
Tunnel local IP address     : 10.1.1.1
Retail service              : N/A
SSID                        : "1"
Previous Access Point IP    : (Not Specified)
IMSI                        : 206100000000001
MGW router                  : 300
Mobile Gateway              : 10.0.0.29
APN                         : full.dotted.apn.mnc010.mcc206.gprs
Last move time              : 2014/03/24 15:38:52
-------------------------------------------------------------------------------
No. of UE: 1

*A:Dut-C# show router 50 wlan-gw soft-gre-tunnels local-ip 10.1.1.1 remote-
ip 239.0.0.2 isa-group 1
===============================================================================
```

```
Soft GRE tunnels
===============================================================================
Remote IP address           : 239.0.0.2
Local IP address            : 10.1.1.1
ISA group ID                : 1
ISA group member ID         : 5
Time established            : 2014/03/24 15:38:52
Number of UE                : 1
Access Point MAC            : 00:00:00:00:00:01
AP MAC learn failed         : false

Tunnel QoS
----------
Operational state           : active
Number of UE                : 1
Remaining hold time (s)     : N/A
-------------------------------------------------------------------------------
No. of tunnels: 1

*A:Dut-C# show router 50 wlan-gw soft-gre-tunnels local-ip 10.1.1.1 remote-
ip 239.0.0.2 isa-group 1 member 5
===============================================================================
Soft GRE tunnels
===============================================================================
Remote IP address           : 239.0.0.2
Local IP address            : 10.1.1.1
ISA group ID                : 1
ISA group member ID         : 5
Time established            : 2014/03/24 15:38:52
Number of UE                : 1
Access Point MAC            : 00:00:00:00:00:01
AP MAC learn failed         : false

Tunnel QoS
----------
Operational state           : active
Number of UE                : 1
Remaining hold time (s)     : N/A
-------------------------------------------------------------------------------
No. of tunnels: 1

*A:Dut-C# show router 50 wlan-gw soft-gre-tunnels local-ip 10.1.1.1 remote-
ip 239.0.0.2 isa-group 1 member 5 summary
===============================================================================
Soft GRE tunnels summary
===============================================================================
Remote IP address - Local IP address
-------------------------------------------------------------------------------
239.0.0.2 - 10.1.1.1
-------------------------------------------------------------------------------
No. of tunnels: 1


*A:Dut-C# show router 50 wlan-gw soft-gre-tunnels local-ip 10.1.1.1 remote-
ip 239.0.0.2 isa-group 1 member 5 detail
===============================================================================
Soft GRE tunnels
===============================================================================
Remote IP address           : 239.0.0.2
```

```
Local IP address            : 10.1.1.1
ISA group ID                : 1
ISA group member ID         : 5
Time established            : 2014/03/24 15:38:52
Number of UE                : 1
Access Point MAC            : 00:00:00:00:00:01
AP MAC learn failed         : false

Tunnel QoS
----------
Operational state           : active
Number of UE                : 1
Remaining hold time (s)     : N/A
Service Access Points(SAP)
===============================================================================
Service Id         : 2147483650
SAP                : 5/1/lo-gre:1             Encap             : q-tag
Description        : Internal SAP
Admin State        : Up                       Oper State        : Up
Flags              : None
Multi Svc Site     : None
Last Status Change : 03/24/2014 15:03:48
Last Mgmt Change   : 03/24/2014 15:14:00
-------------------------------------------------------------------------------
Encap Group Specifics
-------------------------------------------------------------------------------
Encap Group Name   : _tmnx_SHAPER_GR000       Group Type        : ISID
Qos-per-member     : TRUE
Members            :
1
-------------------------------------------------------------------------------
QOS
-------------------------------------------------------------------------------
E. qos-policy      : 1                        Q Frame-Based Acct: Disabled
E. Sched Policy    :                          E. Agg-limit      : -1
                                              Limit Unused BW   : Disabled
-------------------------------------------------------------------------------
Encap Group Member 1 Base Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A

Forwarding Engine Stats
                        Packets                 Octets
For. InProf           : 0                       0
For. OutProf          : 0                       0
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
-------------------------------------------------------------------------------
Encap Group Member 1 Queue Statistics
-------------------------------------------------------------------------------
Packets                 Octets

Egress Queue 1
For. InProf           : 0                       0
For. OutProf          : 0                       0
Dro. InProf           : 0                       0
Dro. OutProf          : 0                       0
===============================================================================
-------------------------------------------------------------------------------
```

```
No. of tunnels: 1
```

tunnels

Syntax **tunnels** [**local-ip***ip-address*] [**remote-ip** *ip-address*] [**isa-group** *wlan-gw-group-id*] [**member** [1 to 255]] [**summary**] [**detail**]
tunnels local-ip *ip-address* **remote-ip** *ip-address* **ue**

Context show>router>wlan-gw

Description This command displays tunnel operation information.

Parameters *ip-address* — Specifies the local IP address of this system that is the destination IP address in the tunnel header of received packets

Values

ip-address:	ipv4-address - a.b.c.d
ipv6-address :	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D

ip-address — Specifies the remote IP address of the Mobile Gateway that is the source IP address in the tunnel header of received packets

Values

ip-address:	ipv4-address - a.b.c.d
ipv6-address :	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D

wlan-gw-group-id — Specifies the identifier of the WLAN gateway ISA group that terminates GRE for this group interface

Values 1 to 4

member — Specifies the identifier of this WLAN gateway ISA group membe.

Values 1 to 255

summary — Displays a summary of the specified parameter.

detail — Displays detailed information

ue — Displays information for the specified user equipment

Output The following output is an example of WLAN-GW tunnel information.

Sample Output

Note: The remote/local IP addresses are locally generated for VLAN tunnels.

```
show router 50 wlan-gw tunnels
===============================================================================
Access Point tunnels
===============================================================================
Remote IP address           : fe80::3e8f:ffff:fe00:1901
Local IP address            : fe80::ff:fe02:202
ISA group ID                : 1
ISA group member ID         : 4
Time established            : 2015/01/07 17:42:01
Number of UE                : 1
Access Point MAC            : 00:00:00:00:00:05
AP MAC learn failed         : false
Encapsulation               : vlan
VLAN tag 1                  : 1000
VLAN tag 2                  : (None)
-------------------------------------------------------------------------------
No. of tunnels: 1
===============================================================================
```

radius-server-policy

Syntax **radius-server-policy** *policy-name* [**acct-on-off**]
radius-server-policy *policy-name* **associations**
radius-server-policy ***policy-name*** **msg-buffer-stats**
radius-server-policy *policy-name* **statistics**
radius-server-policy [**acct-on-off**]

Context show>aaa

Description This command displays RADIUS server policy information.

Parameters *policy-name* — Displays information pertaining to the specified policy name

associations — Displays the association between the RADIUS server policy and the applications referencing the policy (RADIUS proxy, route downloader, authentication policy, accounting policy, dynamic services policy)

statistics — Displays statistics of the RADIUS server policy and RADIUS servers referenced in the policy

acct-on-off — Displays the acct-on-off operational state for the RADIUS server policy

msg-buffer-stats — Displays statistics for the RADIUS message buffering

Output The following is an example of RADIUS server policy information.

Sample Output

```
show aaa radius-server-policy "aaa-server-policy-1"
===============================================================================
RADIUS server policy "aaa-server-policy-1"
===============================================================================
Description                 : Radius AAA server policy
Acct Request script policy  : (Not Specified)
Auth Request script policy  : (Not Specified)
Accept script policy        : script-policy-1
Acct-On-Off                 : Enabled (state Not Blocked)
-------------------------------------------------------------------------------
RADIUS server settings
-------------------------------------------------------------------------------
Router                      : "Base"
Source address              : (Not Specified)
Access algorithm            : direct
Retry                       : 3
Timeout (s)                 : 5
Hold down time (s)          : 30
Last management change      : 02/20/2013 13:32:05
===============================================================================
===============================================================================
Servers for "aaa-server-policy-1"
===============================================================================
Idx Name                            Address         Port        Oper State
                                                    Auth/Acct
-------------------------------------------------------------------------------
1   server-1                        172.16.1.1      1812/1813   in-service
===============================================================================


# show aaa radius-server-policy acct-on-off
===============================================================================
RADIUS server policies AcctOnOff state
===============================================================================
Name                                  OperState      LastStateChange
-------------------------------------------------------------------------------
aaa-server-policy-1                   on             02/20/2013 21:23:57
aaa-server-policy-2                   NotApplicable  NotApplicable
aaa-server-policy-3                   sendAcctOn     NotApplicable
aaa-server-policy-4                   off            02/20/2013 21:40:57
-------------------------------------------------------------------------------
No. of policies: 4
===============================================================================

# show aaa radius-server-policy "aaa-server-policy-1" acct-on-off
===============================================================================
RADIUS server policy "aaa-server-policy-1" AcctOnOff info
===============================================================================
Oper state                  : on
Session Id                  : 242FFF0000000451253EED
Last state change           : 02/20/2013 21:23:57
Trigger                     : startUp
Server                      : "server-1"
===============================================================================


show aaa radius-server-policy "aaa-server-policy-3" msg-buffer-
```

```
stats
===============================================================================
RADIUS server policy "aaa-server-policy-3" message buffering stats
===============================================================================
buffering acct-interim      : enabled
  min interval (s)          : 60
  max interval (s)          : 3600
  lifetime (hrs)            : 12
buffering acct-stop         : enabled
  min interval (s)          : 60
  max interval (s)          : 3600
  lifetime (hrs)            : 12

Statistics
-------------------------------------------------------------------------------
Total acct-stop messages in buffer                        : 6
Total acct-interim messages in buffer                     : 10
Total acct-stop messages dropped (lifetime expired)       : 0
Total acct-interim messages dropped (lifetime expired)    : 0
Last buffer clear time                                    : N/A
Last buffer statistics clear time                         : N/A
-------------------------------------------------------------------------------
===============================================================================


show aaa radius-server-policy "aaa-server-policy-1" statistics
===============================================================================
RADIUS server policy "aaa-server-policy-1" statistics
===============================================================================
Tx transaction requests                         : 383
Rx transaction responses                        : 383
Transaction requests timed out                  : 0
Transaction requests send failed                : 0
Packet retries                                  : 0
Transaction requests send rejected              : 0
Authentication requests failed                  : 0
Accounting requests failed                      : 0
Ratio of access-reject over auth responses      : 0%
Transaction success ratio                       : 100%
Transaction failure ratio                       : 0%
Statistics last reset at                        : n/a

Server 1 "server-1" address 172.16.1.1 auth-port 1812 acct-port 1813
-------------------------------------------------------------------------------
Tx request packets                              : 383
Rx response packets                             : 383
Request packets timed out                       : 0
Request packets send failed                     : 0
Request packets send failed (overload)          : 0
Request packets waiting for reply               : 0
Response packets with invalid authenticator     : 0
Response packets with invalid msg authenticator : 0
Authentication packets failed                   : 0
Accounting packets failed                       : 0
Avg auth response delay (10 100 1K 10K) in ms   :   27.1   22.8   22.8   22.8
Avg acct response delay (10 100 1K 10K) in ms   :   6.24   12.5   11.5   11.5
Statistics last reset at                        : n/a

===============================================================================
```

```
show aaa radius-server-policy "myRadiusServerPolicy1" associations
===============================================================================
RADIUS Proxy Associations
===============================================================================
Router RADIUS Proxy Server Purpose Username
-------------------------------------------------------------------------------
Base myProxyServerBase acc (default)
vprn10 myProxyServer1 acc (default)
-------------------------------------------------------------------------------
No. of associations: 2


show aaa radius-server-policy "aaa-server-policy-1" associations
===============================================================================
RADIUS Proxy Associations
===============================================================================
Router RADIUS Proxy Server Purpose Username
-------------------------------------------------------------------------------
Base   myProxyServerBase   acc     (default)
-------------------------------------------------------------------------------
No. of associations: 1
===============================================================================
No route downloader entries found.
===============================================================================
Authentication Policy Associations
===============================================================================
Authentication Policy
-------------------------------------------------------------------------------
auth-policy-1
-------------------------------------------------------------------------------
No. of associations: 1
===============================================================================
===============================================================================
Accounting Policy Associations
===============================================================================
Accounting Policy
-------------------------------------------------------------------------------
acct-policy-1
acct-policy-2
-------------------------------------------------------------------------------
No. of associations: 2
===============================================================================
No dynamic-services policy entries found.
```

wlan-gw-group

Syntax **wlan-gw-group** *wlan-gw-group-id*
wlan-gw-group *wlan-gw-group-id* **associations**
wlan-gw-group *wlan-gw-group-id* **member** *member-id*
wlan-gw-group *wlan-gw-group-id* **member** *member-id* **resource-statistics**
wlan-gw-group *wlan-gw-group-id* **member** *member-id* **statistics** [**type** *type*] [**non-zero-value-only**]

wlan-gw-group

Context show>isa

Description This command displays WLAN-GW group information, including WLAN-GW tunnels.

Parameters *wlan-gw-group-id* — Displays information about the specified WLAN-GW group ID

Values 1 to 4

associations — Displays information about associations for the specified WLAN-GW group ID

member *member-id* — Displays information about the WLAN-GW-specific status and basic statistics information about the specified member

Values 1 to 255

type *type* — Displays a reduced output to only show statistics of the specified type

Values packet-errors, host-errors, bd-errors, forwarding, reassembly, aa, radius, arp, dhcp, dhcp6, icmp, icmp6

non-zero-value-only — Displays a reduced output to only show statistics whose value is bigger than zero

resource-statistics — Displays the resource usage on the specified group member

statistics — Displays statistics information about the members of the specified WLAN-GW group

Output The following output is an example of ISA WLAN-GW group information.

Sample Output

```
*A:Dut-B>config>isa>wlan-gw-group$ show isa wlan-gw-group 4
===============================================================================
WLAN Gateway group 4
===============================================================================
Administrative state        : in-service
Operational state           : in-service
Degraded                    : false
Active IOM limit            : 0
Active MDA limit            : 14
Port policy                 : (Not Specified)
Tunnel port policy          : (Not Specified)
Dsm ISA AA group            : (Not Specified)
Last Mgmt Change            : 06/28/2017 15:07:34
-------------------------------------------------------------------------------
NAT specific information for ISA group 4
-------------------------------------------------------------------------------
Reserved sessions           : 0
High Watermark (%)          : (Not Specified)
Low Watermark (%)           : (Not Specified)
Accounting policy           : (Not Specified)
UPnP mapping limit          : 524288
Suppress LsnSubBlksFree     : false
LSN support                 : enabled
Last Mgmt Change            : 06/28/2017 15:06:40
```

```
-------------------------------------------------------------------------------
===============================================================================
===============================================================================
ISA Group 4 members
===============================================================================
Group Member     State          Mda  Addresses  Blocks     Se-% Hi Se-Prio
-------------------------------------------------------------------------------
4     1          active         3/1  0          0          < 1  N  0
4     2          active         3/2  0          0          < 1  N  0
4     3          active         4/1  0          0          < 1  N  0
4     4          active         4/2  0          0          < 1  N  0
4     5          active         5/1  0          0          < 1  N  0
4     6          active         5/2  0          0          < 1  N  0
4     7          active         6/1  0          0          < 1  N  0
4     8          active         6/2  0          0          < 1  N  0
4     9          active         7/1  0          0          < 1  N  0
4     10         active         7/2  0          0          < 1  N  0
4     11         active         8/1  0          0          < 1  N  0
4     12         active         8/2  0          0          < 1  N  0
4     13         active         9/1  0          0          < 1  N  0
4     14         active         9/2  0          0          < 1  N  0
-------------------------------------------------------------------------------
No. of members: 14
```

isa-filter

Syntax **isa-filter**
isa-filter *name*
isa-filter *name* **associations**
isa-filter *name* **ipv4**
isa-filter *name* **ipv6**

Context show>subscr-mgm

Description This command displays ISA filter information.

Parameters *name* — Specifies the ISA filter name, up to 32 characters

associations — Displays associated information about the specified ISA filter name

ipv4 — Display IPv4 ISA filter information for the specified ISA filter name

ipv6 — Display IPv6 ISA filter information for the specified ISA filter name

isa-policer

Syntax **isa-policer** *policer-name*
isa-policer *policer-name* **associations**
isa-policer

Context show>subscr-mgm

Description This command displays ISA policer information.

Parameters *policer-name* — Specifies the ISA policer name, up to 32 characters

associations — Displays associated information about the specified ISA policer name

gtp-session

Syntax **gtp-session imsi** *imsi* **apn** *apn-string* | **gtp-session** [**mgw-address** *ip-address*] [**mgw-router** *router-instance*] [**remote-control-teid** *teid*] [**local-control-teid** *teid*] [**detail**]
gtp-session imsi *imsi*
gtp-statistics

Context show>subscr-mgmt>wlan-gw

Description This command displays GTP session information

Parameters *imsi* — Specifies the IMSI (International Mobile Subscriber Identity) of this UE

apn-string — Specifies the APN (Access Point Name)

ip-address — Specifies the IP address of the Mobile Gateway, that is the source IP address in the tunnel header of received packets

router-instance — Specifies the identifier of the virtual router instance where the GTP tunnel is terminated

teid — Specifies the remote control plane Tunnel Endpoint Identifier (TEID)

teid — Specifies the local control plane TEID

detail — Displays detailed information

Output The following is an example of subscriber management WLAN-GW GTP session information.

Sample Output

```
*A:Dut-C# show subscriber-mgmt wlan-gw gtp-session
===============================================================================
GTP sessions
===============================================================================
IMSI                        : 206100000000001
APN                         : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router       : 300
Mobile Gateway address      : 10.20.1.2
Remote control TEID         : 66592
Local control TEID          : 4291821824
Charging characteristics    : (None)
Uplink AMBR (kbps)          : (None)
Downlink AMBR (kbps)        : (None)
IMSI                        : 206100000000002
APN                         : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router       : 300
```

```
Mobile Gateway address      : 10.20.1.2
Remote control TEID         : 66608
Local control TEID          : 4290773248
Charging characteristics    : (None)
Uplink AMBR (kbps)          : (None)
Downlink AMBR (kbps)        : (None)
IMSI                        : 206100000000033
APN                         : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router       : "Base"
Mobile Gateway address      : 10.20.1.2
Remote control TEID         : 66560
Local control TEID          : 4293918976
Charging characteristics    : (None)
Uplink AMBR (kbps)          : (None)
Downlink AMBR (kbps)        : (None)
IMSI                        : 206100000000034
APN                         : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router       : "Base"
Mobile Gateway address      : 10.20.1.2
Remote control TEID         : 66576
Local control TEID          : 4292870400
Charging characteristics    : (None)
Uplink AMBR (kbps)          : (None)
Downlink AMBR (kbps)        : (None)
-------------------------------------------------------------------------------
No. of GTP sessions: 4
===============================================================================
*A:Dut-C#

*A:Dut-C# show subscriber-mgmt wlan-gw gtp-
session imsi 206100000000001 apn full.dotted.apn.mnc010.mcc206.gprs
===============================================================================
GTP session
===============================================================================
IMSI                        : 206100000000001
APN                         : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router       : 300
Mobile Gateway address      : 10.20.1.2
Remote control TEID         : 66592
Local control TEID          : 4291821824
Charging characteristics    : (None)
Uplink AMBR (kbps)          : (None)
Downlink AMBR (kbps)        : (None)
Bearer 5
  rem TEID                  : 1073808421
  loc TEID                  : 4291821861
  uplink GBR (kbps)         : 0
  uplink MBR (kbps)         : 4992
  downlink GBR (kbps)       : 0
  downlink MBR (kbps)       : 1984
  QoS Class ID              : 8
  alloc/ret priority        : 1
===============================================================================
*A:Dut-C#

*A:Dut-C# show subscriber-mgmt wlan-gw gtp-session imsi 206100000000001
```

```
===============================================================================
GTP sessions
===============================================================================
IMSI                        : 206100000000001
APN                         : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router       : 300
Mobile Gateway address      : 10.20.12
Remote control TEID         : 66592
Local control TEID          : 4291821824
Charging characteristics    : (None)
Uplink AMBR (kbps)          : (None)
Downlink AMBR (kbps)        : (None)
Bearer 5
  rem TEID                  : 1073808421
  loc TEID                  : 4291821861
  uplink GBR (kbps)         : 0
  uplink MBR (kbps)         : 4992
  downlink GBR (kbps)       : 0
  downlink MBR (kbps)       : 1984
  QoS Class ID              : 8
  alloc/ret priority        : 1
-------------------------------------------------------------------------------
No. of GTP sessions: 1
===============================================================================
*A:Dut-C#

*A:Dut-C# show subscriber-mgmt wlan-gw gtp-session mgw-address 10.20.12
===============================================================================
GTP sessions
===============================================================================
IMSI                        : 206100000000001
APN                         : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router       : 300
Mobile Gateway address      : 10.20.12
Remote control TEID         : 66592
Local control TEID          : 4291821824
Charging characteristics    : (None)
Uplink AMBR (kbps)          : (None)
Downlink AMBR (kbps)        : (None)
IMSI                        : 206100000000002
APN                         : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router       : 300
Mobile Gateway address      : 10.20.12
Remote control TEID         : 66608
Local control TEID          : 4290773248
Charging characteristics    : (None)
Uplink AMBR (kbps)          : (None)
Downlink AMBR (kbps)        : (None)
IMSI                        : 206100000000033
APN                         : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router       : "Base"
Mobile Gateway address      : 10.20.12
Remote control TEID         : 66560
Local control TEID          : 4293918976
Charging characteristics    : (None)
```

```
Uplink AMBR (kbps)          : (None)
Downlink AMBR (kbps)        : (None)
IMSI                        : 206100000000034
APN                         : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router       : "Base"
Mobile Gateway address      : 10.20.12
Remote control TEID         : 66576
Local control TEID          : 4292870400
Charging characteristics    : (None)
Uplink AMBR (kbps)          : (None)
Downlink AMBR (kbps)        : (None)
-------------------------------------------------------------------------------
No. of GTP sessions: 4
===============================================================================
*A:Dut-C#

*A:Dut-C# show subscriber-mgmt wlan-gw gtp-session mgw-address 10.20.12 mgw-
router "Base"
===============================================================================
GTP sessions
===============================================================================
IMSI                        : 206100000000033
APN                         : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router       : "Base"
Mobile Gateway address      : 10.20.12
Remote control TEID         : 66560
Local control TEID          : 4293918976
Charging characteristics    : (None)
Uplink AMBR (kbps)          : (None)
Downlink AMBR (kbps)        : (None)
IMSI                        : 206100000000034
APN                         : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router       : "Base"
Mobile Gateway address      : 10.20.12
Remote control TEID         : 66576
Local control TEID          : 4292870400
Charging characteristics    : (None)
Uplink AMBR (kbps)          : (None)
Downlink AMBR (kbps)        : (None)
-------------------------------------------------------------------------------
No. of GTP sessions: 2
===============================================================================
*A:Dut-C#

*A:Dut-C# show subscriber-mgmt wlan-gw gtp-session mgw-router 300
===============================================================================
GTP sessions
===============================================================================
IMSI                        : 206100000000001
APN                         : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router       : 300
Mobile Gateway address      : 10.20.12
Remote control TEID         : 66592
Local control TEID          : 4291821824
Charging characteristics    : (None)
```

```
Uplink AMBR (kbps)          : (None)
Downlink AMBR (kbps)        : (None)
IMSI                        : 206100000000002
APN                         : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router       : 300
Mobile Gateway address      : 10.20.12
Remote control TEID         : 66608
Local control TEID          : 4290773248
Charging characteristics    : (None)
Uplink AMBR (kbps)          : (None)
Downlink AMBR (kbps)        : (None)
-------------------------------------------------------------------------------
No. of GTP sessions: 2
===============================================================================
*A:Dut-C#

*A:Dut-C# show subscriber-mgmt wlan-gw gtp-session remote-control-teid 66560
===============================================================================
GTP sessions
===============================================================================
IMSI                        : 206100000000033
APN                         : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router       : "Base"
Mobile Gateway address      : 10.20.12
Remote control TEID         : 66560
Local control TEID          : 4293918976
Charging characteristics    : (None)
Uplink AMBR (kbps)          : (None)
Downlink AMBR (kbps)        : (None)
-------------------------------------------------------------------------------
No. of GTP sessions: 1
===============================================================================
*A:Dut-C# show subscriber-mgmt wlan-gw gtp-session local-control-teid 4292870400
===============================================================================
GTP sessions
===============================================================================
IMSI                        : 206100000000034
APN                         : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router       : "Base"
Mobile Gateway address      : 10.20.12
Remote control TEID         : 66576
Local control TEID          : 4292870400
Charging characteristics    : (None)
Uplink AMBR (kbps)          : (None)
Downlink AMBR (kbps)        : (None)
-------------------------------------------------------------------------------
No. of GTP sessions: 1
===============================================================================
*A:Dut-C#

*A:Dut-C# show subscriber-mgmt wlan-gw gtp-session mgw-address 10.20.12 mgw-
router "Base" local-control-teid 4292870400 remote-control-teid 66576
===============================================================================
GTP sessions
===============================================================================
IMSI                        : 206100000000034
```

```
APN                         : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router       : "Base"
Mobile Gateway address      : 10.20.12
Remote control TEID         : 66576
Local control TEID          : 4292870400
Charging characteristics    : (None)
Uplink AMBR (kbps)          : (None)
Downlink AMBR (kbps)        : (None)
-------------------------------------------------------------------------------
No. of GTP sessions: 1
===============================================================================

*A:Dut-C# show subscriber-mgmt wlan-gw gtp-session mgw-address 10.20.12 mgw-
router "Base" local-control-teid 4292870400 remote-control-teid 66576 detail
===============================================================================
GTP sessions
===============================================================================
IMSI                        : 206100000000034
APN                         : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router       : "Base"
Mobile Gateway address      : 10.20.12
Remote control TEID         : 66576
Local control TEID          : 4292870400
Charging characteristics    : (None)
Uplink AMBR (kbps)          : (None)
Downlink AMBR (kbps)        : (None)
Bearer 5
  rem TEID                  : 1073808405
  loc TEID                  : 4292870437
  uplink GBR (kbps)         : 0
  uplink MBR (kbps)         : 4992
  downlink GBR (kbps)       : 0
  downlink MBR (kbps)       : 1984
  QoS Class ID              : 8
  alloc/ret priority        : 1
-------------------------------------------------------------------------------
No. of GTP sessions: 1
===============================================================================
*A:Dut-C#
```

gtp-statistics

Syntax **gtp-statistics**

Context show>subscr-mgmt>wlan-gw

Description This command displays GTP statistics.

Output The following is an example of WLAN-GW GTP statistics.

Sample Output

```
*A:Dut-C# show subscriber-mgmt wlan-gw gtp-statistics
```

```
===============================================================================
GTP statistics
===============================================================================
tx echo requests                                        : 33
tx echo responses                                       : 0
tx errors                                               : 0
rx echo requests                                        : 0
rx echo responses                                       : 33
rx errors                                               : 0
rx version not supported                                : 0
rx zero TEID responses                                  : 0
path faults                                             : 0
path restarts                                           : 0
tx invalid msgs                                         : 0
tx create PDP context requests                          : 4
tx create PDP context responses                         : 0
tx delete PDP context requests                          : 0
tx delete PDP context responses                         : 0
tx create session requests                              : 0
tx create session responses                             : 0
tx delete session requests                              : 0
tx delete session responses                             : 0
tx delete bearer requests                               : 0
tx delete bearer responses                              : 0
tx create bearer responses                              : 0
tx update bearer responses                              : 0
tx modify bearer requests                               : 0
tx modify bearer responses                              : 0
tx error indication count                               : 0
rx invalid msgs                                         : 0
rx create PDP context requests                          : 0
rx create PDP context responses                         : 4
rx delete PDP context requests                          : 0
rx delete PDP context responses                         : 0
rx create session requests                              : 0
rx create session responses                             : 0
rx delete session requests                              : 0
rx delete session responses                             : 0
rx delete bearer requests                               : 0
rx delete bearer responses                              : 0
rx create bearer requests                               : 0
rx update bearer requests                               : 0
rx modify bearer requests                               : 0
rx modify bearer responses                              : 0
rx error indication count                               : 0
rx invalid pkt length                                   : 0
rx unknown pkts                                         : 0
rx missing IE pkts                                      : 0
rx bad IP header pkts                                   : 0
rx bad UDP header pkts                                  : 0
rx discarded pkts                                       : 0
rx in-session discarded pkts                            : 0
rx pkts                                                 : 37
tx discarded pkts                                       : 0
tx pkts                                                 : 37
===============================================================================
*A:Dut-C#
```

ssid

Syntax **ssid**

Context show>subscr-mgmt>wlan-gw

Description This command displays SSID information.

statistics

Syntax **statistics**

Context show>subscr-mgmt>wlan-gw

Description This command displays statistics information.

tunnels

Syntax **tunnels** [**router** *router-name*] [**remote-ip** *ip-address*] [**local-ip** *ip-address*] [**encapsulation** *encap* [*encap*...(**upto 3 max**)]] [**qtag1** *qtag*] [**qtag2** *qtag*] [**ap-sap** *sap-id*] [**min-num-ue** *minimum*] [**max-num-ue** *maximum*] [**ap-mac-learn-failed** {**true** | **false**}] [**get-num-results**] [**addr-family** *family*] [**ue-type** *ue-type* [*ue-type*...(**upto 5 max**)]]

Context show>subscr-mgmt>wlan-gw

Description This command displays all the WLAN-GW tunnels matching the specified criteria. Unlike the similar command in the **show**>**router**>**vprn** context, this command also includes information on tunnels containing ISA-only UEs such as migrant, DSM and l2-wholesale.

Parameters *router-name* — The name or ID of the router where the tunnel terminates

ip-address — The IPv4 or IPv6 address indicating one, or both, of the tunnel endpoint IP addresses

encap — The tunnel encapsulation type, for example GRE, L2TP, or VLAN

qtag — The Q-tags specifying the l2-ap-delimiting tags

sap-id — The SAP-ID of the l2-ap SAP

minimum — The minimum number of UEs on the tunnel, after applying the UE type filter

maximum — The maximum number of UEs on the tunnel, after applying the UE type filter

ap-mac-learn-failed true | false — Filters the results to display only tunnels that have learned the AP-MAC (**false**) or have not learned the AP-MAC (**true**).

get-num-results — Displays the total number of tunnels at the end of each tunnel record

family — Specifies the tunnel's IP family type (IPv4 or IPv6)

ue-type — Filters the display based on the presence of specified UE types and is used in conjunction with **min-num-ue** and **max-num-ue**

Values migrant, dsm, l2w, esm, or xcon

Output The following is an example of WLAN-GW tunnels.

Sample Output

```
Node# show subscriber-mgmt wlan-gw tunnels
===============================================================================
Access Point tunnels
===============================================================================
Router                       : 50
Encapsulation                : gre
Remote IP address            : 192.0.2.1
Local IP address             : 192.0.2.2
-------------------------------------------------------------------------------
First move time              : N/A
ISA group ID                 : 1
ISA member ID                : 3
Interface                    : grp-vprn_ue-2/1/2:50
Interface Service ID         : 4
AP MAC address               : 00:53:00:00:00:05
AP MAC learn failed          : false
AP SAP                       : (Unknown)
Remote UDP port              : N/A
Tag 1                        : N/A
Tag 2                        : N/A
No. of UE                    : 1
No. of migrant UE            : 0
No. of DSM UE                : 1
No. of layer-2 wholesale UE  : 0
No. of cross-connect UE      : 0
No. of ESM UE                : 0
-------------------------------------------------------------------------------
No. of Access point tunnels: 1
===============================================================================
```

ue

Syntax **ue** [**vlan** *qtag*] [**mpls-label** *label*] [**retail-svc-id** *service-id*] [**ssid** *service-set-id*] [**previous-access-point** *ip-address*] [**bd** *bridge-id*]

ue mac *ieee-address* [**bd** *bridge-id*]

Context show>subscr-mgmt>wlan-gw

Description This command displays user equipment (UE) information.

Parameters *qtag* — Displays information about the VLAN Q-tag present in the traffic received from this UE.

Values 1 to 4095

label — Displays information about the MPLS label present in the traffic received from this UE

service-id — Specifies an existing service ID. If no *svc-id* is specified then it indicates that the interface is a network interface in the Base router instance

Values {*id* | *svc-name*}

id: 1 to 2147483647

svc-name: Specifies an existing service name, up to 64 characters (*svc-name* is an alias for input only. The *svc-name* gets replaced with an id automatically by SR OS in the configuration)

service-set-id — Displays information about the Service Set ID (SSID) of this UE

ip-address — Displays information about the IP address of the previous Access Point (AP) of this UE

bridge-id — Displays specified HLE bridge domain information of this UE

Values 1 to 4294967295

ieee-address — Displays information about the MAC address of this UE

Values xx:xx:xx:xx:xx:xx or xx-xx-xx-xx-xx-xx

Output The following displays WLAN-GW information.

Sample Output

```
System# show subscriber-mgmt wlan-gw ue
===============================================================================
User Equipments
===============================================================================
MAC address                 : 00:02:00:00:00:39
-------------------------------------------------------------------------------
VLAN Q-tag                  : 1
MPLS label                  : (Not Specified)
Tunnel router               : 50
Tunnel remote IP address    : 20C9::7:1:2
Tunnel local IP address     : 2032::1:1:7
Retail service              : N/A
SSID                        : 1
Previous Access Point IP    : (Not Specified)
IMSI                        : (Not Specified)
Last move time              : 2013/07/02 07:45:31

-------------------------------------------------------------------------------
No. of UE: 1
===============================================================================
System#
```

11.25.2.8.2 Debug Commands

call-trace

Syntax **call-trace**

Context debug

Description This command enables the context to set up various call-trace debug sessions.

wlan-gw

Syntax [**no**] **wlan-gw**

Context debug

Description This node contains all the parameters to set up specific call-trace debug sessions for WLAN-GW.

The **no** form of this command will stop all configured WLAN-GW traces.

statistic

Syntax **statistic type** *type* **name** *name*
no statistic

Context debug>wlan-gw>group

Description This command enables debugging of the specified statistic. The first packet that causes an increase of the specified statistic is shown in debug output. After the first packet, debugging of the counter is stopped.

Parameters *type* — Displays the type of statistic to be debugged; for example, DHCP or RADIUS

Values packet-errors, host-errors, bd-errors, forwarding, reassembly, aa, radius, arp, dhcp, dhcp6, icmp, icmp6

name — Specifies the name, up to 256 characters, of the statistic within that group. For a complete list, see the command **show isa wlan-gw-group** *wlan-gw-group-id* **member** *member-id* **statistics**.

ue

Syntax **ue** *ieee-address* [**profile** *trace-profile-name*]
no ue *ieee-address*

Context debug>call-trace>wlan-gw

Description This command starts tracing the UE with the specified MAC address. The trace is started with default parameters or optionally parameters specified in the trace-profile.

The **no** form of this command stops the trace and make sure no new traces are started.

Parameters *ieee-address* — Displays information about the MAC address of this UE

trace-profile-name — Specifies the name of a configured trace profile

11.25.2.8.3 Tools Commands

acct-on

Syntax **acct-on** [**radius-server-policy** *policy-name*] [**force**]

Context tools>perform>aaa

Description This command triggers a RADIUS Accounting-On message:

- for all radius-server-policies that have acct-on-off configured.
- for the specified radius-server-policy if the acct-on-off is configured

The Accounting-On message is not sent when the last successful event for the RADIUS server policy was an Accounting-On message. In this case, an Accounting-Off should be sent first. By specifying the keyword **force**, this is overruled.

Parameters *policy-name* — Specifies the **radius-server-policy** for which the Accounting-On should be sent

force — Sends an Accounting-On also if the last successful event was an Accounting-On

acct-off

Syntax **acct-off** [**radius-server-policy** *policy-name*] [**force**] [**acct-terminate-cause** *number*]

Context tools>perform>aaa

Description This command triggers a RADIUS Accounting-Off message:

- for all radius-server-policies that have acct-on-off configured
- for the specified radius-server-policy if the acct-on-off is configured

The Accounting-Off message is not sent when the last successful event for the radius server policy was an Accounting-Off message. In this case, an Accounting-On should be sent first. By specifying the keyword **force**, this is overruled.

Parameters *policy-name* — Specifies the **radius-server-policy** for which the Accounting-Off should be sent

force — Sends an Accounting-On also if the last successful event was an Accounting-Off

number — Overrides the default Acct-Terminate-Cause (User-Request) in the Accounting-Off message

radius-acct-terminate-cause

Syntax **radius-acct-terminate-cause**

Context tools>dump>aaa

Description This command shows all available termination causes and their respective number values. The TermCause is equivalent to VSA 226 **alc-error-code** numeric values. The description is equivalent to VSA 227**alc-error-message** string.

radius-server-policy

Syntax **radius-server-policy** *policy-name* **msg-buffer** [**session-id** *acct-session-id*]

Context tools>perform>aaa
tools>dump>aaa

Description This command dumps the RADIUS message buffer content for the specified radius-server-policy:

- message-type (acct-interim or acct-stop)
- Acct-Session-Id
- Remaining lifetime

When specifying the session-id, the message details are displayed.

Parameters *policy-name* — Specifies the **radius-server-policy** for which the message buffer content should be displayed

acct-session-id — Displays the RADIUS message details for the message with specified session-id that is stored in the RADIUS message buffer

performance

Syntax **performance mda** *mda-id* **last** *time-span time-unit*

Context tools>dump>wlan-gw>isa

Description This command generates an overview of the processing load and data processed by the specified ISA over a period of time. The following time periods are supported:

- last minute with seconds granularity
- last hour with minutes granularity
- last day with hours granularity
- last day with days granularity

Parameters *mda-id* — Specifies the MDA for getting performance measurements in *slot*/*mda* format

Values slot — 1 to 10
mda — 1 to 2

time-span — Specifies the period for which to get measurements

Values 1 to 60 (sec), 1 to 60 (min), 1 to 24 (hrs), 1 (days)

time-unit — Specifies the period for which to get measurements

Values sec, min, hrs, days

Output This command displays performance information.

Sample Output

```
Node# /tools dump wlan-gw isa performance mda 2/1 last 5 min
===============================================================================
Measurements for last 5 minutes on Slot #2 MDA #1
===============================================================================
Timestamp            |     Wait     Idle     Work | Total jobs |    Total data
---------------------+----------------------------+------------+---------------
01/22/2018 10:14:04  |   99.47%    0.53%    0.00% |          0 |          - -
01/22/2018 10:13:41  |   99.46%    0.54%    0.00% |          3 |          3 Kb
01/22/2018 10:12:41  |   99.47%    0.53%    0.00% |          0 |          - -
01/22/2018 10:11:41  |   99.47%    0.53%    0.00% |          0 |          - -
01/22/2018 10:10:41  |   99.45%    0.55%    0.00% |          0 |          - -
===============================================================================
```

ue

Syntax **ue** [**wlan-gw-group** *wlan-gw-group-id*] [**mda** *mda-id*] [**next-index** *index*] [**summary**] [**detail**] [**bd** *bridge-id*] [**ue-mac** *ieee-address*] [**ue-vlan** *vlan*] [**state-description** *state*] [**tunnel-router** *router-instance*] [**tunnel-source-ip** *ip-address*] [**tunnel-destination-ip** *ip-address*] [**tunnel-type** *tunnel-type*] [**ue-ip** *ipv4-address*] [**dhcp6-addr** *ipv6-address*] [**slaac-prefix** *ipv6-address*] [**aggregate-summary**]

Context tools>dump>wlan-gw

Description This command dumps user equipment (UE) information.

The **summary** option displays a count of UEs per ISA and the **aggregate-summary** displays a count of matched UEs over the whole WLAN-GW.

Output This command displays UE information.

Sample Output

```
tools dump wlan-gw ue
===============================================================================
Matched 1 session on Slot #4 MDA #1
===============================================================================
UE-Mac          : 00:02:00:00:00:11     UE-vlan         : 3600
UE IP Addr      : N/A                   UE timeout      : N/A
UE IP6 Addr     : N/A
Description     : L2-user
Auth/CoA-time   : 01/07/2015 18:56:01
Tunnel MDA      : 5/1                   Tunnel Router   : 50
MPLS label      : N/A                   Shaper          : Default
Tunnel Src IP   : 203.0.113.235         Tunnel Dst IP   : 10.1.1.1
Tunnel Type     : GRE
Anchor SAP      : 4/1/nat-out-ip:2049.6
AP-Mac          : Unknown               AP-RSSI         : Unknown
AP-SSID         : Unknown
Last-forward    : 01/07/2015 18:56:01   Last-move       : None
Session Timeout : None                  Idle Timeout    : 300 sec
Acct Update     : None                  Acct Interval   : N/A
Acct Session-Id : N/A
Acct Policy     : N/A
NAT Policy      : N/A
Redirect Policy : N/A
IP Filter       : N/A
App-profile     : N/A
Rx Oper PIR     : N/A                   Rx Oper CIR     : N/A
Tx Oper PIR     : N/A                   Tx Oper CIR     : N/A
Rx Frames       : 0                     Rx Octets       : 0
Tx Frames       : 0                     Tx Octets       : 0
-------------------------------------------------------------------------------
===============================================================================
No sessions on Slot #4 MDA #2 match the query
No sessions on Slot #5 MDA #1 match the query
No sessions on Slot #5 MDA #2 match the query
```

11.25.2.8.4 Clear Commands

radius-server-policy

Syntax **radius-server-policy** *policy-name* **msg-buffer** [**acct-session-id** *acct-session-id*]

radius-server-policy *policy-name* **statistics** [**msg-buffer-only**]
radius-server-policy *policy-name* **server** *server-index* **statistics**

Context clear>aaa

Description This command dumps the RADIUS message buffer content for the specified radius-server-policy:

- message-type (acct-interim or acct-stop)
- Acct-Session-Id
- Remaining lifetime

When specifying the session-id, the message details are displayed.

Parameters *policy-name* — Specifies the radius-server-policy for which the information should be cleared

acct-session-id — Deletes all RADIUS messages or the RADIUS message with specified session-id from the RADIUS message buffer

msg-buffer-only — Clears all statistics for the specified radius-server-policy: radius-server-policy statistics, RADIUS server statistics and RADIUS message buffer statistics. With the optional keyword “msg-buffer-only”, only the RADIUS message buffer statistics are cleared

server-index — Clears the RADIUS server statistics for the specified server-index in the specified radius-server-policy

isa-subnets

Syntax **isa-subnets all**
isa-subnets interface *ip-int-name*
isa-subnets prefix *ipv6-address/prefix-length*

Context clear>router>wlan-gw

Description This command clears specific subnets from the pool-manager. Associated UE’s is removed from the system.

When clearing the last subnet on an ISA the pool-manager will automatically allocate a new subnet with allocation-level 0%.

Parameters **all** — Clears all the isa-subnets

ip-int-name — Clears all the isa-subnets of a specific subscriber-interface

ipv6-address/prefix-length — Clears a specific IPv6 address and prefix length

wlan-gw-group

Syntax **wlan-gw-group** *group-id* **member** *member-id* **resource-peak-values**
wlan-gw-group *group-id* **member** *member-id* **statistics**

Context clear>wlan-gw>isa

Description This command resets wlan-gw statistics per group member.

Parameters *group-id* — Specifies the WLAN-GW group ID

Values 1 to 4

member *member-id* — Specifies the member ID

Values 1 to 255

statistics — Resets the statistics measurements to zero

resource-peak-values — Resets the resource peak values to the current resource measurements

12 GTP

12.1 GTP Uplink

SR OS supports subscriber traffic forwarding over an uplink GTP tunnel towards a GGSN or P-GW. This requires a per-subscriber GTP tunnel based on authentication to be configured. Each subscriber may access only a single APN. Both GTPv1 (Gn) and GTPv2 (S2a/S2b) are supported. A single primary PDP context per subscriber is supported on the Gn interface (3GPP TS 29.060 Release 8) from SR OS to the GGSN. A single default-bearer per subscriber is supported on the S2b interface (3GPPTS 29.274 Release 10), and S2a interface (3GPP TS 29.274 Release 11) from SR OS to the P-GW.

12.1.1 Identification Attributes

GTP requires at least an IMSI and an APN to set up a connection. The IMSI is required to identify the user, and the APN is required to identify the network the user is connecting to.

There is a 1:1 relationship between IMSI and subscriber ID. It is possible to provision only one of these and the other will be accepted as the same value. If both are provisioned, they must be equal. Therefore, it is not possible to set up more than one GTP tunnel per subscriber.

APN can be provisioned explicitly per subscriber, or a default APN can be provisioned per VRF. If this APN does not contain an Operator Identifier (OI), it will be added automatically based on the IMSI.

12.1.2 P-GW/GGSN Selection

The initial address of the P-GW or GGSN can either be provided during authentication, or, in absence of authentication, resolved dynamically via DNS. For DNS, a FQDN is generated based on the APN as specified in 3GPP 29.303. This FQDN always consists of both the Network-ID (NI) portion and the OI portion, and is formatted as NI.apn.epc.mnc<MNC>.mcc<MCC>.3gppnetwork.org. The system will perform an S-NAPTR lookup with this FQDN.

When multiple GW addresses are returned as part of this lookup, load-balancing is performed according to regular NAPTR and SRV rules. If no addresses are returned, or the S-NAPTR lookup failed, the system tries a regular A host lookup with the same FQDN. In this case, SR OS load-balances over multiple GW addresses using a round-robin mechanism. DNS servers can be configured per VPRN, except for the base router where the servers defined in the BOF are used.

```
config>router
config>service>vprn
    apn "internet.mno1.apn"
config>service>vprn>dns
    primary-dns 10.1.1.1
    secondary-dns 10.1.1.2
    tertiary-dns 10.1.1.3
    ipv4-source-address 10.1.1.1
exit
```

After initial GTP setup, it is possible for the P-GW or GGSN to return another address as a GTP-C or GTP-U destination. All data plane traffic is sent using the signaled GTP-U address. All subsequent control plane traffic is forwarded to the new GTP-C address.

The VRF used for GTP tunneling can be selected via provisioning a retain service ID for the subscriber. The source IP for GTP tunneling is taken from a loopback interface with the name **system** in that VRF. If no such interface is present, the tunnel setup fails.

12.1.3 Configuration

Profiles with signaling-related configuration per mobile gateway can be created locally on the SR OS router. These profiles include configuration for the interface type used between the router and the mobile gateway, path management parameters, retransmission parameters, and default values for GTP information elements such as AMBR. Each profile can be mapped to a specific GW IP address or subnet per VRF. Most of the per-session/context parameters can be overridden via RADIUS authentication. Refer to the *7750 SR and VSR RADIUS Attributes Reference Guide* for more information.

```
config>subscr-mgmt>gtp>peer-profile
    no description
    interface-type s2a
    ip-ttl 255
    keep-alive interval 60 retry-count 4 timeout 5
    message-retransmit timeout 5 retry-count 3
    protocol-configuration-options apco
    no python-policy
    rat-type wlan
    no report-wlan-location
```

```
        session-hold-time 30
        charging-characteristics
            no home
            no roaming
        exit
        ggsn
            qos
                no ambr
                arp 1
                down-link gbr 2000 mbr 2000
                up-link gbr 5000 mbr 5000
            exit
        exit
        pgw
            qos
                ambr down-link 20000 up-link 10000
                arp 1
                down-link gbr 0 mbr 0
                qci 8
                up-link gbr 0 mbr 0
            exit
        exit

    config>router
    config>service>vprn
        gtp uplink peer-profile-map
            address 198.51.100.1/32 "pgw-policy"
            address 239.0.113.0/24 "ggsn-policy"
        exit
```

12.1.4 QoS Support

SR OS provides appropriate traffic treatment and remarking based on DSCP bits in the outer and inner header in GTP packets.

Downstream from PGW/GGSN, the DSCP bits from the outer header in a GTP packet can be mapped to a forwarding class on network ingress, which can be preserved through the chassis as the packet passes to the egress IOM. On egress, reclassification can be done based on either the inner or outer DSCP bits, depending on the configuration value of the **use-ingress-l2tp-dscp** option in the SLA profile.

In the upstream direction, regular ESM FC classification is used. This FC is carried through the IOM to the egress complex. In the egress complex, this FC can be used for remarking of the outer DSCP values.

DSCP and default FC values for egress GTP-C packets can be configured under **sgt-qos**.

It is possible to signal the subscriber's aggregate rate or the rate of a specific scheduler in the downlink AMBR IE in both GTPv1-C and GTPv2-C. This uses the **report-rate** configuration of the SLA profile; the pppoe-actual-rate and rfc5515-actual-rate values are not applicable for GTP. This value can be subtracted with a value signaled during authentication to take into account an average use for selective breakout.

```
config>subscr-mgmt>sla-prof>egress
    report-rate
```

Other signed QoS IE values are taken from static configuration or values signaled in authentication.

12.1.5 GTP Session Hold

Deletion of an IPoE or PPPoE session also triggers deletion of any corresponding GTP sessions. This deletion is subject to a configurable hold time. When the subscriber returns with the same GGSN/P-GW parameters within the hold time, the GTP connection will not be re-signaled. This avoids releasing resources (such as IP addresses) too quickly on the GGSN/P-GW. This is useful in the following cases:

- Non-seamless control plane mobility with WLAN-GW. Due to SHCV usage, the UE will be completely removed, but should reconnect quickly, preferably with the same IP address. Having a short hold time avoids the P-GW releasing the IP address.
- IP handover from WiFi-GW to RAN. In some cases, WiFi connectivity is lost before wireless data connectivity is established. By holding the GTP session active on the WLAN-GW, the IP address is also held on the GGSN/P-GW until the RAN connection is complete.

While a GTP session is in hold, all downstream traffic is dropped, but no error indication messages are sent.

12.1.6 Selective Breakout

This feature adds support for selecting a subset of traffic from a host (via IP filter) for local forwarding, while tunneling the remaining traffic to GGSN/PGW. This allows selected traffic to bypass the mobile packet core. The IP address for the host still comes from the GGSN/PGW during GTP session setup. Therefore, the selected traffic for local breakout from SR OS requires NAT functionality in order to draw the return traffic back to the router. To support address overlap within GTP, the NAT functionality is L2-aware. The selection of traffic for local breakout (local forwarding and NAT) is based on a net action in an upstream Ip filter applied to the host.

Selective breakout can be enabled on a per-host basis via RADIUS VSA (Alc-GTP-Local-Breakout) in access-accept. It is not possible to change this during a host's lifetime, such as via CoA. AA functionality is supported for local breakout traffic. Also, LI (after NAT) is supported for local breakout traffic, and is enabled via existing secure CLI, as stated in the *OAM and Diagnostics Guide*.

```
system>config>filter
    ip-filter 10 create
        entry 1 create
            match protocol udp
                dst-port eq 4000
            exit
            action gtp-local-breakout
        exit
```

On traffic ingress from the host UE, normal ESM host lookup and CAM lookup with the ingress host filter is performed. If there is a match in the filter indicating "gtp-local-breakout", the traffic is forwarded within the chassis to an ISA-BB, where is it subjected to L2-aware NAT function, and afterwards is forwarded using regular routing in the NAT outside VRF. The inside IP address is the address returned in GTP, and may not match a NAT L2-aware inside prefix. The outside IP is an address belonging to the NAT outside IP address range on the ISA. If the filter action results in a "forward" action (default), the traffic is GTP-tunneled without performing NAT functionality. The traffic received from the network can be a normal L3 packet or a GTP encapsulated packet. The normal Layer 3 packet is expected to be destined to the NAT outside IP and is normally routed to the NAT ISA.

By default, per-host accounting includes counters that are aggregated across GTP and local breakout traffic. Separate counters can be obtained by directing the GTP and local breakout traffic into different queues associated with the corresponding ESM host based on QoS IP classification. NAT information (outside IP and port range) associated with an ESM host subjected to selective breakout is included in accounting-updates.

Selective breakout is supported for IPv4 only.

12.1.7 IPoE Support

A GTPv2 session or GTPv1 PDP context is set up when IPoE session authentication includes any GTP parameters. The GTP session provides the IPv4 and IPv6 address used for the connecting host. Currently only DHCPv4 and SLAAC are supported to deliver these addresses back to the client. If DHCP is used, SR OS automatically derives a standards-compliant subnet mask and default GW from the signaled IP address. It is important that all GTP subscribers are in a shared split-horizon domain and that there is no Layer 2 switching between GTP subscribers. Only a single IPoE session is supported per GTP subscriber. Additionally, DNS and NBNS can be signaled via GTP (A)PCO and reflected in DHCP, SLAAC, and stateless DHCPv6. Control plane packets such as DHCP and ICMPv6 RS are always terminated on the BNG and are not forwarded over GTP.

Figure 189 shows a sample IPoE session for GTP.

Figure 189 **IPoE Session**

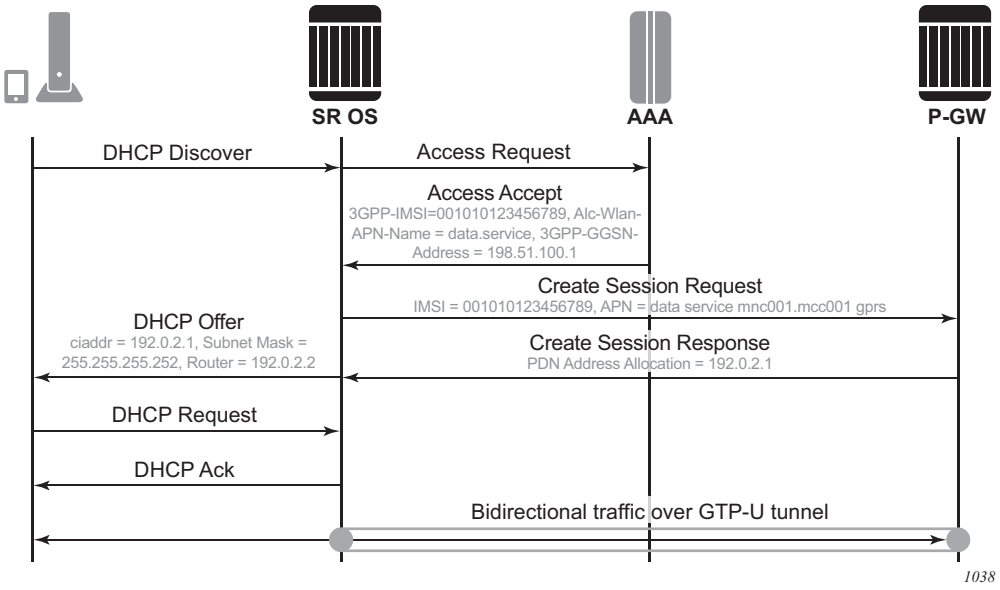

GTP without an IPoE session is available for IPv4 DHCP leases only for backwards compatibility. It may not be used for new deployments; existing deployments should move to the IPoE session concept.

12.1.8 PPPoE Support

A GTPv2 session or GTPv1 PDP context is set up when PPPoE session authentication includes any GTP parameters. The GTP session provides the IPv4 and IPv6 address to be used for the connecting host. IPCP and IPv6CP with SLAAC are supported to signal these addresses to the client. Only a single PPPoE session is supported per GTP subscriber. Additionally, DNS and NBNS can be signaled via GTP (A)PCO and reflected in IPCP, SLAAC, and stateless DHCPv6. Control plane packets such as ICMPv6 RS are always terminated on the BNG and are not forwarded over GTP.

12.2 GTP Access

SR OS supports TPSDA functionality over GTP access funnels towards eNodeBs using the S11 (GTPv2-C) and S1-U (GTP-U) interfaces toward MME and eNodeB. In this setup, the SR OS router performs the roles of BNG, SGW, and PGW. Both IPv4 and IPv6 connectivity over GTP is supported. GTP session authentication can be performed using LUDB, RADIUS, or NASREQ. GTP access termination is built on top of forwarding path extensions.

Figure 190 **GTP Access**

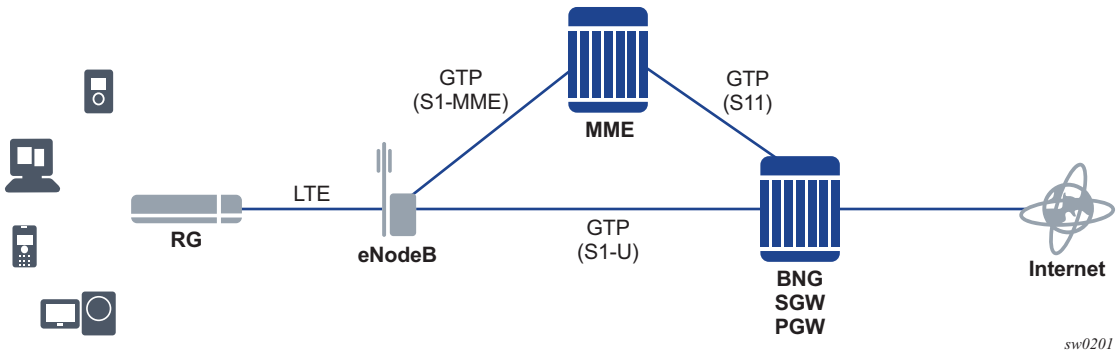

12.2.1 GTP Termination

GTP S11 termination can be enabled on interfaces in the base router and VPRNs. The configuration is linked to an APN policy that lists all supported APNs and the authentication mechanism (for example, LUDB or RADIUS) to be used per APN. The configured APN string should match the value signaled in GTP; however, fallback configuration is supported for any unknown APNs.

```
*A:FWGW>config>subscr-mgmt>gtp# info
----------------------------------------------
            apn-policy apn_demo create
                apn demo.mnc001.mcc001.gprs create
                    radius-auth-policy "gtp_auth"
                exit
            exit
----------------------------------------------
*A:FWGW>config>service>vprn# info
----------------------------------------------
            gtp
                s11
                    interface gtp_endpoint create
                        apn-policy "apn_demo"
                    exit
                exit
```

```
            exit
            interface "gtp_endpoint" create
                description "Tunnel endpoint IP"
                address 192.0.2.10/32
                loopback
            exit
-----------------------------------------------
```

A GTP peer profile defines certain signaling parameters such as TTL values, keepalive timers, retransmit timers, and default values for certain information elements. By default, an automatically-created profile with the name “default_s11” is used. A more specific profile can be configured and mapped to a peer by a per-VRT mapping of IP address or prefix to that profile.To map all peers within the same VRF to the same profile, it is possible to use prefix 0.0.0.0/0.

```
*A:FWGW>config>subscr-mgmt>gtp# info
-----------------------------------------------
            peer-profile "s11_peers" create
                interface-type s11
                keep-alive interval 180 retry-count 10 timeout 20
                message-retransmit timeout 3 retry-count 1
                mme
                    qos
                        ambr down-link 50000 up-link 10000
                    exit
                exit
            exit
-----------------------------------------------
*A:FWGW>config>service>vprn>gtp>s11# info
-----------------------------------------------
                    peer-profile-map
                        address 0.0.0.0/0 peer-profile s11_peers
                    exit
-----------------------------------------------
```

When an S11 session is set up, the accompanying S1-U bearer is terminated in the same VRF, but it is directly linked to a group interface in either the same or a different VRF. A default group interface can be configured per APN, which can be overridden during S11 session authentication. The used group interfaces are of type **gtp** and require an FPE construct to operate. The traffic takes two passes through the forwarding plane. For upstream data, the GTP-U header is stripped in the first pass and inserted into the group interface for regular ESM processing, based on existing IPoE functionality, in the second pass. For downstream data, regular ESM processing is performed in the first pass, then the traffic egresses over the group interface and GTP-U encapsulation occurs in the second pass.

12.2.1.1 Multiple APNs

SR OS supports multiple APN connectivity for the same IMSI. When an additional session setup with a different APN is triggered, the system will treat this as a completely new session that goes through APN selection and authentication again. PDN session user plane can be instantiated in the same VRF as previous sessions or in different VRFs. Sessions for the same IMSI can belong to the same ESM subscriber by using the same subscriber-id, but are not required to do so. S11 messages defined with UE granularity in 3GPP TS 29.274 will apply to all sessions of the same IMSI.

12.2.2 GTP Session Setup

When a GTPv2 Create Session Request message arrives, SR OS looks up the required authentication mechanism in the applicable APN policy. To aid in the identification process, both RADIUS and NASREQ can be configured with GTP specific include attributes such as IMSI, MSISDN, and IMEI. If a PAP message is present in the PCO IE of the Create Session request, that username and password is used for authentication; otherwise the system falls back to IMSI as the username and password configured in the policy.

For LUDB-based authentication, it is recommended to use **derived-id** for identification values.

Authentication is performed per GTP session and not per IP stack (host). Therefore, the initial authentication returns parameters for all stacks that need to be set up.

After a successful authentication, a Create Session Response message is sent, which includes all relevant parameters including assigned addresses, DNS servers, and applicable QoS values. The Create Session Response message is followed by an initial Modify Bearer Request message. When the host setup is completed with a Modify Bearer Response message, downstream data can then flow towards the eNodeB. If IPv6 is enabled, an unsolicited RA is sent.

Figure 191 shows a high-level overview of the setup call flow using RADIUS authentication.

Figure 191 **High-level Example of GTP Access Setup**

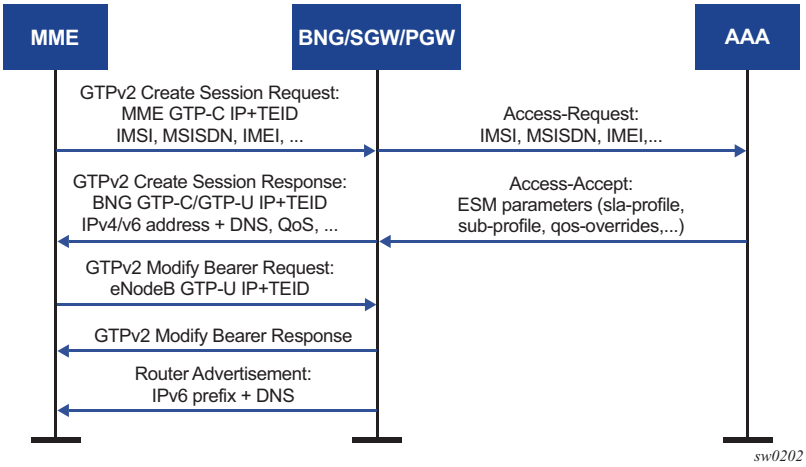

12.2.2.1 Supported IP Stacks

IPv4 and IPv6 SLAAC are supported via dual-stack bearers. IP addresses can be provisioned during authentication or through Local Address Assignment. Any DNS (IPv4 and IPv6) and NBNS (IPv4 only) addresses are signaled in the PCO IE. An unsolicited RA is sent after the first Modify Bearer Request message is received. An RS-triggered RA is also supported and either RA can be configured to contain DNS servers. Stateless DHCPv6 can also be used to retrieve the DNS configuration.

12.2.3 QoS

ESM QoS using subscriber and SLA profiles also applies to GTP hosts. To align with 3GPP signaling, it is possible to enable the ingress or egress **report-rate** in the SLA profile. When specified, the signaled UL/DL APN-AMBR values are taken from the referenced report rate object (that is, the scheduler or subscriber aggregate rate). If this mechanism is not used, these values can be taken from AAA signaling (VSA 3GPP-GPRS-Negotiated-QoS-Profile in RADIUS). As a last resort, default values can be found in the peer profile. Similarly, QCI and ARP values can be retrieved from AAA and have default values in the peer profile.

SR OS maintains the FC classification and in- and out-of-profile markings consistent over the two dataplane passes. The egress and ingress QoS policies in the SLA profile should be configured with the following considerations.

- Enable **de-mark** for access egress and map each FC to the dot1p as defined in Table 166, thus ensuring that the GTP-U encapsulated packet uses the same classification as the ESM context.
- Perform classification for access ingress based on dot1p as defined in Table 166 and enable **in-profile** and **de-1-out-of-profile** for each FC, thus ensuring that the ESM context uses the same classification as determined for the GTP-U packet. However, a different classification scheme can be used, if required; for example, based on DSCP or IP criteria.

Table 166 **FC to dot1p Mapping**

FC	dot1p
be	0
l2	1
af	2
l1	3
h2	4
ef	5
h1	6
nc	7

12.2.4 Multicast

GTP supports the reception of MLD and IGMP, with either a redirect-interface or per-host replication enabled. Per-SAP replication is not applicable and is ignored.

12.3 DHCP over GTP-u

A routed gateway (RG) with a fixed-wireless WAN link can get an IPv4 address and an IPv6 /128 address using DHCPv4 and DHCPv6 IA_NA, respectively. The RG can also get an IPv6 prefix using DHCPv6 PD for it's LAN. A fixed-wireless RG can be an RG with integrated LTE or 5G modem, or connected to an external 5G modem using Ethernet or WLAN. Once the default bearer and PDU session is created using NAS signaling, DHCP messages can be forwarded by the RG (with integrated modem), or the standalone modem (referred to as user equipment (UE)), over the default bearer towards the BNG, which functions as fixed-wireless gateway and terminates the GTP tunnel. DHCP messages are received and sent by the BNG over the GTP tunnel.

12.3.1 Address Management Related PCOs

Without DHCPv4, the fixed-wireless RG can only get an IPv4 address during initial attach using NAS SM (session management) procedures. As part of NAS signaling, the BNG acting as a fixed-wireless gateway allocates an IPv4 address and returns it in a PDN address allocation (PAA) IE in the GTP session creation response.

- If an RG requires IPv4 address allocation using DHCPv4 instead of address allocation during initial attach using NAS signaling, then it can signal a PCO "deferred address allocation" (value 0x000b) in GTP session setup. If this PCO is present in GTP setup request, then the BNG will not allocate an IPv4 address during GTP session setup. The GTP setup will complete with IPv4 address of 0.0.0.0 returned in a PAA to the RG. The BNG will only allocate and assign an IPv4 address using subsequent DHCPv4 from the RG, which is received over the GTP-u tunnel (S1-U interface) on the BNG.
- If PCO "allocation via NAS" (value 0x000a) is present in a GTP setup request, then the IPv4 address is allocated during GTP setup and returned in a GTP session creation response.
- If neither PCO is present, then by default, anIPv4 address is allocated on the GTP setup, and returned in the GTP session creation response. This can be overridden using the **skip-gtp-ipv4-alloc** CLI command. In this case, no IPv4 address is allocated during GTP setup and is only allocated using a subsequent DHCPv4 exchange.

The **skip-gtp-ipv4-alloc** configuration is a per-APN configuration. This can be overridden on a per-subscriber basis using a RADIUS VSA (`Alc-Gtp-Skip-Ipv4-Alloc-Override`).

12.3.2 Address Allocation Modes

DHCPv6 IA_NA can be used for IPv6 WAN address assignment. IPv6 WAN prefix can also be supported using SLAAC.

For DHCPv4, the following address allocation modes are supported.

- DHCPv4 proxy. The IPv4 address offered in DHCPv4 from the RG can be provided from AAA during initial authentication in the framed-IP-address RADIUS attribute.
- DHCPv4 relay to a local or external DHCPv4 server for IPv4 address allocation. The pool can be provided from AAA during initial authentication in the framed-pool RADIUS attribute, and can be added in the relayed DHCP packets with a DHCP Option 82 vendor-specific sub-option.

For DHCPv6, the following address allocation modes are supported.

- DHCPv6 proxy. The /128 IPv6 address for RG WAN using DHCPv6 IA_NA can be provided from AAA during initial authentication in Alc-IPv6-address VSA. The PD prefix requested by RG using DHCPv6 IA_PD can be provided from AAA in the Delegated-IPv6-Prefix RADIUS attribute.
- DHCPv6 relay to a local or external DHCPv6 server for IPv6 address and prefix allocation. The pool for /128 IPv6 address allocation for DHCPv6 IA_NA can be provided by AAA during initial authentication in the framed-ipv6-pool attribute. Pool for PD prefix can be provided by AAA during initial authentication in the Alc-Delegated-IPv6-Pool VSA. These are signaled by DHCPv6 relay to the server.

Authentication is performed only during GTP session setup. All the RADIUS attributes related to DHCP and IP stack (host) setup are received during this authentication. No subsequent authentication is performed on the receiving DHCP. A single IPoE session containing all the hosts (DHCPv4, DHCPv6 IA_NA, and SLAAC) are created per GTP session. PD prefix should be configured to be added as a managed-route if both DHCPv6 IA_NA and SLAAC are in use and creating respective hosts.

Figure 192 shows the call flow for GTP session setup, subsequent DHCP exchange, and host creation in case of RG with a standalone modem (user equipment is noted as UE in the diagram).

Figure 192 **Call Flow for GTP Session Setup**

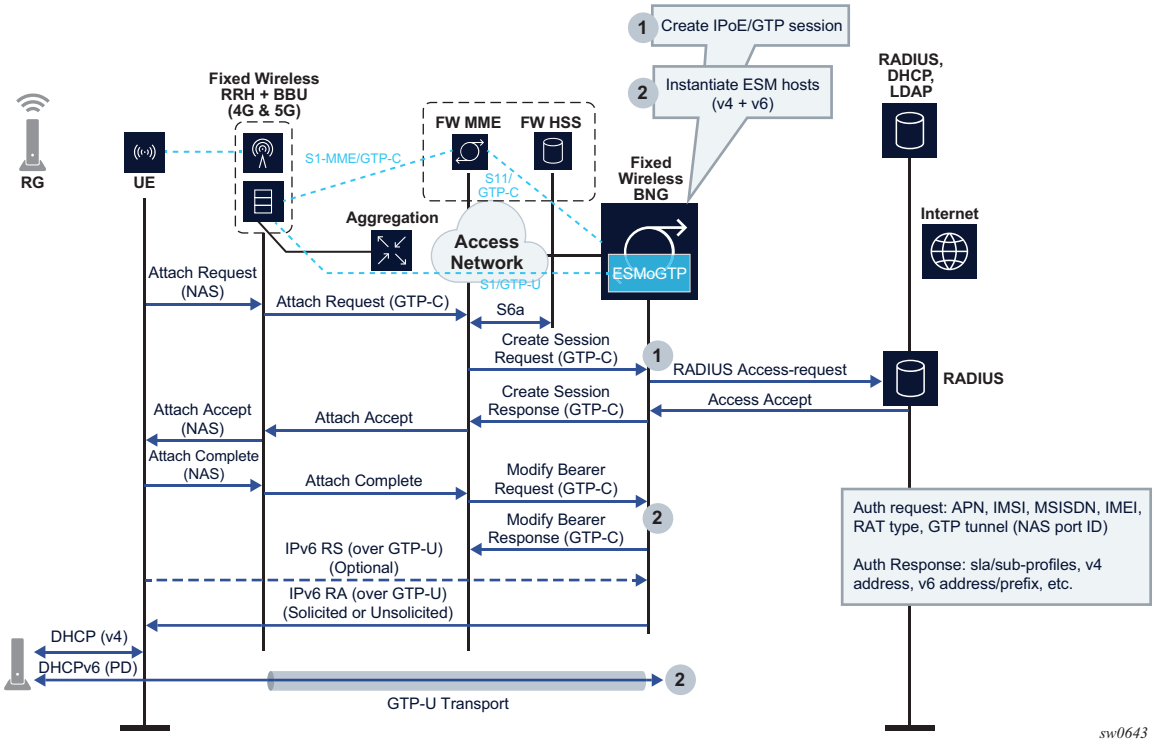

DHCP over GTP-u tunnel is not supported in conjunction with bonding (hybrid-access).

DHCP over GTP-u tunnel is supported in conjunction with multiple APNs. The APNs can have different address management (for example, in case of two APNs). It is possible to use WAN address assignment using NAS for one APN and DHCPv4 for another APN.

Broadcast flag should not be set in DHCP messages from the client.

12.4 GTP Peering

SR OS tracks each GTP-C peer for which it has at least a single GTP session or PDP context active. It tracks the peer's operational state with the following mechanisms:

- Regular GTP echo messages and parameters are configurable on a per-mgw-profile basis. When the echo mechanism fails, the peer is considered down.
- Active route entries towards the peer are monitored. If no route towards the peer is available, the peer is considered down.
- The Restart counter value of the peer is monitored. This is initially learned when the first active session or context is created. If the value is not available in regular messaging, an echo request is sent out immediately to learn the correct value. If the Restart counter is incremented during any later messaging exchange, the peer is considered rebooted.

When a peer is considered down or rebooted, all active GTP sessions and PDP contexts are forcefully removed.

SR OS also keeps a recovery counter in a persistent state, and increments this value on every reboot. This value is kept in the restcntr.txt file on CF3 and may not be modified or removed. This value is included in every control plane message.

SR OS responds to GTP echo messages for both active peers and unknown sources. This can be restricted using CPM filters if required. An incoming echo request from an unknown source will not create a peer; this can only be done by setting up GTP sessions or PDP contexts.

12.5 GTP Command Reference

12.5.1 Command Hierarchies

- GTP Commands
- GTP Service Commands
- GTP Show Commands
- GTP Clear Commands
- GTP Debug Commands

12.5.1.1 GTP Commands

```
config
    — subscriber-mgmt
        — gtp
            — apn-policy policy-name [create]
            — no apn-policy policy-name
                — apn {apn-name | unknown} [create]
                — no apn {apn-name | unknown}
                    — defaults
                        — group-interface interface-name svc-id service-id
                        — no group-interface
                    — diameter-auth-policy diameter-authentication-policy-name
                    — no diameter-auth-policy
                    — radius-auth-policy radius-authentication-policy-name
                    — no radius-auth-policy
                    — [no] skip-gtp-ipv4-alloc
                    — user-db local-user-database-name
                    — no user-db
            — max-held-sessions max-held-sessions
            — no max-held-sessions
            — peer-profile profile-name [create]
            — no peer-profile profile-name
            — peer-profile profile-name [create]
                — change-reporting-action reporting-action
                — no change-reporting-action
                — charging-characteristics
                    — home bit [bit]
                    — no home
                    — roaming bit [bit.]
                    — no roaming
                — description description-string
                — no description
                — end-marker-count packets
```

— **no end-marker-count**
— **ggsn**
— **qos**
— **ambr down-link** *down-link-kbps* **up-link** *up-link-kbps*
— **no ambr**
— **arp** *arp-value*
— **no arp**
— **down-link gbr** *rate* **mbr** *rate*
— **no down-link**
— **up-link gbr** *rate* **mbr** *rate*
— **no up-link**
— **interface-type** {**gn** | **s2a** | **s2b** | **s11**}
— **no interface-type**
— **ip-ttl** *hops*
— **no ip-ttl**
— **ipv4-mtu** *bytes*
— **no ipv4-mtu**
— **keep-alive** [**interval** *seconds*] [**retry-count** *value*] [**timeout** *retry-seconds*]
— **no keep-alive**
— **message-retransmit** [**timeout** *timeout*] [**retry-count** *value*]
— **no message-retransmit**
— **mme**
— **qos**
— **ambr down-link** *down-link-kbps* **up-link** *up-link-kbps*
— **no ambr**
— **arp** *arp-value*
— **no arp**
— **down-link gbr** *rate* **mbr** *rate*
— **no down-link**
— **qci** *qci-value*
— **no qci**
— **up-link gbr** *rate* **mbr** *rate*
— **no up-link**
— **pgw**
— **qos**
— **ambr down-link** *down-link-kbps* **up-link** *up-link-kbps*
— **no ambr**
— **arp** *arp-value*
— **no arp**
— **down-link gbr** *rate* **mbr** *rate*
— **no down-link**
— **qci** *qci-value*
— **no qci**
— **up-link gbr** *rate* **mbr** *rate*
— **no up-link**
— **protocol-configuration-options** {**apco** | **pco**}
— **no protocol-configuration-options**
— **python-policy** *name*
— **no python-policy**
— **rat-type** {**utran** | **geran** | **wlan** | **gan** | **hspa** | **eutran** | **virtual** | *id*}
— **no rat-type**
— [**no**] **report-wlan-location**
— **session-hold-time remaining-lease-time**

```
                    — session-hold-time seconds
                    — no session-hold-time
                — serving-network mcc mcc-value mnc mnc-value
                — no serving-network

configure
    — router
        — gtp
            — s11
                — interface interface-name [create]
                — no interface interface-name
                    — apn-policy apn-policy-name
                    — no apn-policy
                — peer-profile-map
                    — address ip-prefixlip-prefix-length [peer-profile profile-name]
                    — no address ip-prefixlip-prefix-length
            — [no] uplink
                — apn apn
                — no apn
                — pdn-type {ipv4 | ipv6 | ipv4v6}
                — no pdn-type
                — peer-profile-map
                    — address ip-prefixlip-prefix-length [peer-profile profile-name]
                    — no address ip-prefixlip-prefix-length
```

12.5.1.2 GTP Service Commands

```
configure
    — service
        — vprn service-id
            — gtp
                — s11
                    — interface interface-name [create]
                    — no interface interface-name
                        — apn-policy apn-policy-name
                        — no apn-policy
                    — peer-profile-map
                        — address ip-prefixlip-prefix-length [peer-profile profile-
                            name]
                        — no address ip-prefixlip-prefix-length
                — [no] uplink
                    — apn apn
                    — no apn
                    — pdn-type {ipv4 | ipv6 | ipv4v6}
                    — no pdn-type
                    — peer-profile-map
                        — address ip-prefixlip-prefix-length [peer-profile profile-
                            name]
                        — no address ip-prefixlip-prefix-length
```

12.5.1.3 GTP Show Commands

```
show
    — router
        — gtp
            — peer [peer-profile profile-name] [local-address ip-address] [control protocol]
              [interface-type interface-type]
            — peer remote-address {ip-address | ipv6-address} [udp-port port] [dsm-local-
              ip-address {ip-address | ipv6-address}
            — peer remote-address {ip-address | ipv6-address} [udp-port port] statistics
              [dsm-local-ip-address {ip-address | ipv6-address}
            — s11
                — peer-profile-map
            — uplink
            — uplink mgw-address-cache [arec] [snaptr] [srv]
            — uplink mgw-address-cache apn apn-domain-string
    — subscriber-mgmt
        — gtp
            — peer-profile [profile-name]
            — peer-profile profile-name associations
            — s11
            — s11 session imsi imsi apn apn-string
            — s11 session [peer-address ip-address] [router router-instance] [remote-
              control-teid teid] [local-control-teid teid] [detail]
            — s11 session imsi imsi
            — statistics
            — summary
            — uplink
            — uplink session imsi imsi apn apn-string
            — uplink session [peer-address ip-address] [router router-instance] [remote-
              control-teid teid] [local-control-teid teid] [detail]
            — uplink session imsi imsi
```

12.5.1.4 GTP Clear Commands

```
clear
    — router
        — gtp
            — peer remote-address ip-address [remote-udp-port port]
            — peer remote-address ip-address [remote-udp-port port] statistics
            — uplink
                — mgw-address-cache
    — subscriber-mgmt
        — gtp
            — session all
            — session apn apn
            — session apn apn imsi imsi
            — session imsi imsi
            — statistics
```

12.5.1.5 GTP Debug Commands

```
debug
    — [no] gtp
        — [no] event
        — [no] imsi imsi
        — packet detail-level {high | low}
        — packet mode {all | dropped}
        — no packet
        — [no] peer router router-instance address ip-address [udp-port port]
```

12.5.2 GTP Command Descriptions

12.5.2.1 GTP Commands

12.5.2.1.1 Generic Commands

description

Syntax **description** *description-string*
no description

Context config>subscr-mgmt>gtp>peer-profile

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes any description string from the context.

Default no description

Parameters *description-string* — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

subscriber-mgmt

Syntax **subscriber-mgmt**

Context config

Description This command enables the context to configure subscriber management entities. A subscriber is uniquely identified by a subscriber identification string. Each subscriber can have several DHCP sessions active at any time. Each session is referred to as a subscriber host and is identified by its IP address and MAC address.

All subscriber hosts belonging to the same subscriber are subject to the same hierarchical QoS (HQoS) processing. The HQoS processing is defined in the sub-profile (the subscriber profile). A sub-profile refers to an existing scheduler policy (configured in **the config>qos>scheduler-policy** context) and offers the possibility to overrule the rate of individual schedulers within this policy.

Because all subscriber hosts use the same scheduler policy instance, they must all reside on the same complex.

12.5.2.1.2 GTP Commands

gtp

Syntax **gtp**

Context config>router
config>service>vprn

Description This command enables the context to configure GTP parameters for the routing context.

s11

Syntax **s11**

Context config>router>gtp
config>service>vprn>gtp

Description This command enables GTP configuration related to S11 termination in this VRF.

interface

Syntax **interface** *interface-name* [**create**]
no interface *interface-name*

Context config>router>gtp>s11
config>service>vprn>gtp>s11

Description This command activates GTP termination on the specified interface.

The **no** form of this command disables GTP termination on the specified interface, if there are no active sessions associated with the interface.

Parameters *interface-name* — Specifies the name of the interface. The name must begin with a letter and is up to 32 characters.

create — Creates an entry.

apn-policy

Syntax **apn-policy** *apn-policy-name*
no apn-policy

Context config>router>gtp>s11>interface
config>service>vprn>gtp>s11>interface

Description This command configures an Access Point Name (APN) policy for the S11 interface.

The **no** form of this command removes the APN policy.

Parameters *apn-policy-name* — Specifies the name of the policy, up to 32 characters.

peer-profile-map

Syntax **peer-profile-map**

Context config>router>gtp>s11
config>service>vprn>gtp>s11
config>router>gtp>uplink
config>service>vprn>gtp>uplink

Description This command configures a mapping of addresses and subnets to GTP peer profiles.

address

Syntax **address** *ip-prefix*/*ip-prefix-length* [**peer-profile** *profile-name*]
no address *ip-prefix*/*ip-prefix-length*

Context config>router>gtp>s11>peer-profile-map
config>service>vprn>gtp>s11>peer-profile-map
config>router>gtp>uplink>peer-profile-map
config>service>vprn>gtp>uplink>peer-profile-map

Description This command configures a mapping of an IP address or subnet to a peer profile. If one peer profile is used for the entire router, it is possible to map the entire IPv4 subnet using 0.0.0.0/0.

If no match is found, the default or default S11 peer profile is used.

The **no** form of this command removes the peer profile mapping, affecting only the setup of new peers.

Parameters *ip-prefix*/*ip-prefix-length* — Specifies the IP prefix and prefix length of the subnet.

Values

ipv4-prefix	a.b.c.d (host bits must be 0)
ipv4-prefix-length	[0 to 32]
ipv6-prefix	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D
ipv6-prefix-le	[0 to 128]

profile-name — Specifies the GTP peer profile associated with the address prefix, up to 32 characters.

uplink

Syntax [**no**] **uplink**

Context config>router>gtp
config>service>vprn>gtp

Description This command enables GTP configuration related to a GTP uplink using the Gn, S2a, or S2b interface.

apn

Syntax **apn** *apn*
no apn

Context config>router>gtp>uplink
config>service>vprn>gtp>uplink

Description This command configures the Network Identifier part of the APN.

The **no** form of this command removes the string from the configuration.

Default no apn

Parameters *apn* — Specifies the APN used for this IMSI to connect to this Mobile Gateway, up to 80 characters.

pdn-type

Syntax **pdn-type** {**ipv4** | **ipv6** | **ipv4v6**}
no pdn-type

Context config>router>gtp>uplink
config>service>vprn>gtp>uplink

Description This command configures the PDP type to be signaled in GTP, determining which addresses are requested from the P-GW/GGSN and which hosts are set up afterwards. This can be overridden by RADIUS. If the **ipv4v6** keyword is used, the P-GW/GGSN can fall back to either IPv4 or IPv6.

The **no** form of this command reverts to the default configuration.

Default pdn-typeipv4

Parameters **ipv4** — Specifies the GTP connection requests an IPv4 address.

ipv6 — Specifies the GTP connection requests an IPv6 address.

ipv4v6 — Specifies the GTP connection requests both an IPv4 and an IPv6 address.

gtp

Syntax **gtp**

Context config>subscr-mgmt

Description This command enables the context to configure box-wide GTP parameters and profiles.

apn-policy

Syntax **apn-policy** *policy-name* [**create**]
no apn-policy *policy-name*

Context config>subscr-mgmt>gtp

Description This command configures an APN policy that defines parameters to be used when setting up a new incoming GTP connection. Each APN can be mapped to its own set of parameters.

The **no** form of this command removes the policy from the system. A policy can only be removed if it is not in use.

Parameters *policy-name* — Specifies the name of the policy, up to 32 characters.

create — Creates an entry.

apn

Syntax **apn** {*apn-name* | **unknown**} [**create**]
no apn {*apn-name* | **unknown**}

Context config>subscr-mgmt>gtp>apn-policy

Description This command configures the parameters that should be applied to incoming connections with the APN specified. Multiple APN nodes can be defined per APN policy.

For each APN-policy, one **unknown** APN entry can be created. This APN is used by all connections not matching another APN.

The **no** form of this command removes the APN from the policy. Only new sessions are affected by the removal.

Parameters *apn-name* — Specifies the APN name as it appears in GTP messaging, up to 79 characters.

create — Creates an entry.

defaults

Syntax **defaults**

Context config>subscr-mgmt>gtp>apn-policy>apn

Description This command enables the context to configure default parameters for the GTP connection that can be used when the parameters are not returned in authentication.

group-interface

Syntax **group-interface** *interface-name* **svc-id** *service-id*
no group-interface

Context config>subscr-mgmt>gtp>apn-policy>apn>defaults

Description This command configures the default group interface where the hosts of the GTP connection is enabled. The group interface must be of type **gtp**.

The **no** form of this command removes the default group interface. In this case, a group interface must be specified using authentication.

Default no group-interface

Parameters *interface-name* — Specifies the name of the group interface, up to 32 characters.

service-id — Specifies the ID of the service where the group interface resides.

Values
service-id: 1 to 2147483647
svc-name: up to 64 characters

diameter-auth-policy

Syntax **diameter-auth-policy** *diameter-authentication-policy-name*
no diameter-auth-policy

Context config>subscr-mgmt>gtp>apn-policy>apn

Description This command configures the Diameter authentication policy with which the GTP connection is authenticated.

The **no** form of this command removes the authentication policy. Only new session setups are affected.

Default no diameter-auth-policy

Parameters *diameter-authentication-policy-name* — Specifies the name of the authentication policy to be used, up to 32 characters.

radius-auth-policy

Syntax **radius-auth-policy** *radius-authentication-policy-name*
no radius-auth-policy

Context config>subscr-mgmt>gtp>apn-policy>apn

Description This command configures the RADIUS authentication policy with which the GTP connection is authenticated.

The **no** form of this command removes the authentication policy. Only new session setups are affected.

Default no radius-auth-policy

Parameters *radius-authentication-policy-name* — Specifies the name of the authentication policy to be used, up to 32 characters.

skip-gtp-ipv4-alloc

Syntax [**no**] **skip-gtp-ipv4-alloc**

Context config>subscr-mgmt>gtp>apn-policy>apn

Description This command enables the ability to skip IPv4 address assignment using a GTP session setup response when PCO "allocation via NAS" is not present in a GTP session creation request. Without this configuration, IPv4 address allocation is done using GTP session setup response, even in absence of the PCO "allocation via NAS" in a GTP session setup request.

The **no** form of the command reverts to IPv4 address allocation using GTP.

user-db

Syntax **user-db** *local-user-database-name*
no user-db

Context config>subscr-mgmt>gtp>apn-policy>apn

Description This command configures the LUDB with which the GTP connection is authenticated.

The **no** form of this command removes the user database for authentication with this APN. Only new session setups are affected.

Default no user-db

Parameters *local-user-database-name* — Specifies the name of the LUDB to be used, up to 32 characters.

max-held-sessions

Syntax **max-held-sessions** *max-held-sessions*
no max-held-sessions

Context config>subscr-mgmt>gtp

Description This command configures the maximum number of GTP sessions to be held while their UE is disconnected.

The **no** form of this command reverts to the default.

Default max-held-sessions 2000

Parameters *max-held-sessions* — Specifies the maximum number of GTP sessions.

Values 0 to 262144

peer-profile

Syntax **peer-profile** *profile-name* [**create**]

no peer-profile *profile-name*

Context config>subscr-mgmt>gtp

Description This command creates a new peer profile.

Parameters *profile-name* — Specifies the profile name, up to 32 characters.

create — Creates an entry.

change-reporting-action

Syntax **change-reporting-action** *reporting-action*
no change-reporting-action

Context config>subscr-mgmt>gtp>peer-profile

Description This command specifies the value of the change reporting action IE sends to the peer in applicable messages. The peer needs to indicate support first using the appropriate flag in the indication IE.

This is overridden by AAA, if AAA explicitly request notification changes for either ECGI, TAI or both. If AAA does not request any notification changes or only the generic location change, the configured value is used.

The **no** version of this command indicates that the IE is not sent, unless specified by AAA.

Default no change-reporting-action

Parameters *reporting-action* — Specifies the reporting action value as per TS 29.274.

Values 0 to 255, cgi-sai, rai, tai, ecgi, cgi-sai-rai, tai-ecgi

charging-characteristics

Syntax **charging-characteristics**

Context config>subscr-mgmt>gtp>peer-profile

Description This command enables the context to configure charging characteristics.

Default none

home

Syntax **home** *bit* [*bit*]
no home

Context config>subscr-mgmt>gtp>peer-profile>charging

Description This command configures the charging characteristics for home UE.

Default no home

Parameters *bit* — Specifies up to 16 bits to set in the Charging Characteristics Information Element (IE) for home UE, if not known by other means such as RADIUS.

roaming

Syntax **roaming** *bit* [*bit*]
no roaming

Context config>subscr-mgmt>gtp>peer-profile>charging

Description This command configures the charging characteristics for roaming UE.

The **no** form of this command removes the *bit* value from the configuration.

Default no roaming

Parameters *bit* — Specifies up to 16 bits to set in the Charging Characteristics Information Element (IE) for roaming UE, if not known by other means such as RADIUS.

end-marker-count

Syntax **end-marker-count** *packets*
no end-marker-count

Context config>subscr-mgmt>gtp>peer-profile

Description This command specifies the number of end marker packets that are sent when it is certain no more packets will be sent over the corresponding GTP-U tunnel, such as after a completed mobility event.

The **no** version of this command reverts the value to the default.

Default end-marker-count 1

Parameters *packets* — Specifies the number of end marker packets to send.

Values 0 to 5

ggsn

Syntax **ggsn**

Context config>subscr-mgmt>gtp>peer-profile

Description This command enables the context to configure communication with a GGSN Mobile Gateway.

qos

Syntax **qos**

Context config>subscr-mgmt>gtp>peer-profile>ggsn
config>subscr-mgmt>gtp>peer-profile>mme

Description This command enables the context to configure QoS for a GGSN Mobile Gateway.

mme

Syntax **mme**

Context config>subscr-mgmt>gtp>peer-profile

Description This command configures parameters specific to a Mobility Management Entity (MME) peer.

ambr

Syntax **ambr down-link** *down-link-kbps* **up-link** *up-link-kbps*
no ambr

Context config>subscr-mgmt>gtp>peer-profile>ggsn>qos
config>subscr-mgmt>gtp>peer-profile>pgw>qos
config>subscr-mgmt>gtp>peer-profile>mme>qos

Description This command configures the Aggregated Maximum Bit Rate (AMBR) to be sent in the APN AMBR IE. The contents of this IE can be overridden by RADIUS or report-rate mechanisms. If those mechanisms specify a partial value, such as only specifying the down-link parameter, the other value is picked up from the **ambr** configuration.

For GTPv1, the **no** form of this command implies that the IE will not be sent. If a partial value is received from another source, the missing value will use the following defaults:

- 10000 kb/s up-link
- 20000 kb/s down-link

For GTPv2, the **no** form of this command reverts to the default of 10000 kb/s up-link and 20000 kb/s down-link.

Default no ambr - for ggsn

ambr down-link 20000 up-link 10000 - for mme and pgw

Parameters *down-link-kbps* — Specifies the downlink Aggregated Maximum Bit Rate (AMBR).

Values 0 to 10000000

up-link-kbps — Specifies the uplink AMBR.

Values 0 to 10000000

arp

Syntax **arp** *arp-value*
no arp

Context config>subscr-mgmt>gtp>peer-profile>ggsn>qos
config>subscr-mgmt>gtp>peer-profile>pgw>qos
config>subscr-mgmt>gtp>peer-profile>mme>qos

Description The command configures the allocation/retention priority to be used in the GTP messages as QOS IE (for a Gn interface) or Bearer QOS (for GTPv2).

The **no** form of this command reverts to the default.

Default no ambr - for ggsn

ambr down-link 20000 up-link 10000 - for mme and pgw 1

Parameters *arp-value* — Specifies the Allocation/Retention Priority (ARP).

Values 1 to 3 (for ggsn context)

Values 1 to 15 (for pgw and mme context)

down-link

Syntax **down-link gbr** *rate* **mbr** *rate*
no down-link

Context config>subscr-mgmt>gtp>peer-profile>ggsn>qos
config>subscr-mgmt>gtp>peer-profile>pgw>qos
config>subscr-mgmt>gtp>peer-profile>mme>qos

Description This command configures the down-link bitrate in kb/s to be used in the GTP messages.

The **no** form of this command reverts to the default.

Default down-link gbr 2000 mbr 2000

down-link gbr 2000 mbr 2000 - for ggsn

down-link gbr 0 mbr 0 - for mme and pgw

Parameters **gbr** *rate* — Specifies the downlink Guaranteed Bit Rate (GBR) to be used in the GTP messages as QOS IE (GTPv1) or Bearer QOS (GTPv2).

mbr *rate* — Specifies the downlink Maximum Bit Rate (MBR) to be used in the GTP messages as QOS IE (GTPv1) or Bearer QOS (GTPv2).

up-link

Syntax **up-link gbr** *rate* **mbr** *rate*
no up-link

Context config>subscr-mgmt>gtp>peer-profile>ggsn>qos
config>subscr-mgmt>gtp>peer-profile>pgw>qos
config>subscr-mgmt>gtp>peer-profile>mme>qos

Description This command configures the up-link bitrate in kb/s to be used in the GTP messages.

The **no** form of this command reverts to the default.

Default up-link gbr 5000 mbr 5000

up-link gbr 5000 mbr 5000 - for ggsn

up-link gbr 0 mbr 0 - for mme and pgw

Parameters **gbr** *rate* — Specifies the uplink Guaranteed Bit Rate (GBR) to be used in the GTP messages as QOS IE (GTPv1) or Bearer QOS (GTPv2).

mbr *rate* — Specifies the uplink Maximum Bit Rate (MBR) to be used in the GTP messages as QOS IE (GTPv1) or Bearer QOS (GTPv2).

pgw

Syntax **pgw**

Context config>subscr-mgmt>gtp>peer-profile

Description This command enables the context to configure communication with a Packet Data Network Gateway.

Default none

qci

Syntax **qci** *qci-value*

no qci

Context config>subscr-mgmt>gtp>peer-profile>pgw>qos
config>subscr-mgmt>gtp>peer-profile>mme>qos

Description This command configures the QoS Class Identifier (QCI) to send in the Bearer Level QoS IE in GTPv2 messages.

The **no** form of this command reverts to the default.

Default qci 8

Parameters *qci-value* — Specifies the QCI value to send.

Values 1 to 9

interface-type

Syntax **interface-type** {**gn** | **s2a** | **s2b** | **s11**}
no interface-type

Context config>subscr-mgmt>gtp >peer-profile

Description This command specifies the interface applicable for communications to the peer. If the interface type does not match the given context in an uplink context, the peer setup will fail.

The **no** form of this command reverts to the default value.

Default interface-type s2a

Parameters **gn** — Signaling interface with the peer is Gn as specified in 3GPP TS 29.060.

S2a — Signaling interface with the peer is S2a as specified in 3GPP TS 29.274.

S2b — Signaling interface with the peer is S2b as specified in 3GPP TS 29.274.

S11 — Signaling interface with the peer is S11 as specified in 3GPP TS 29.274.

ip-ttl

Syntax **ip-ttl** *hops*
no ip-ttl

Context config>subscr-mgmt>gtp>peer-profile

Description This command configures the value to put in the IP header's TTL field for GTP control messages.

The **no** form of this command reverts to the default value.

Default ip-ttl 255

Parameters *hops* — Specifies the IP TTL.

Values 1 to 255

ipv4-mtu

Syntax **ipv4-mtu** *bytes*
no ipv4-mtu

Context config>subscr-mgmt>gtp>peer-profile

Description This command configures the value of the IPv4-MTU PCO sent in S11 GTP messages. This is the MTU a device should honor when sending data toward the SGW/PGW. For IPv6, this value is signaled in the RA message and which can be configured in the **group-interface> ipv6>router-advertisements>mtu** context.

The **no** form of this command resets the signaled IPv4 MTU to the default.

Default ipv4-mtu 1400

Parameters *bytes* — Specifies the MTU value in bytes.

Values 512 to 9000

keep-alive

Syntax **keep-alive** [**interval** *seconds*] [**retry-count** *value*] [**timeout** *retry-seconds*]
no keep-alive

Context config>subscr-mgmt>gtp>peer-profile

Description This command configures Echo-Request messages.

The **no** form of this command reverts to the default values.

Default keep-alive interval 60 retry-count 4 timeout 5

Parameters *seconds* — Specifies, in seconds, the interval between keep-alive Echo-Request messages towards the same peer.

Values 0, 60 to 180

Default 60

value — Specifies, in seconds, the interval between keep-alive Echo-Request messages towards the same peer.

Values 1 to 15

Default 4

retry-seconds — Specifies the retry timeout, in seconds.

Values 1 to 20

Default 5

message-retransmit

Syntax **message-retransmit** [**timeout** *timeout*] [**retry-count** *value*]
no message-retransmit

Context config>subscr-mgmt>gtp>peer-profile

Description This command configures the retry-count and response-timeout for GTP messages.

The **no** form of this command reverts to the default values.

Default message-retransmit timeout 5 retry-count 3

Parameters *timeout* — Specifies the interval, in seconds, between retransmission of signaling request messages towards the same peer Mobile Gateway.

Values 1 to 30

Default 5

value — Specifies the number of times a signaling request message is transmitted towards the same peer Mobile Gateway.

Values 1 to 8

Default 3

protocol-configuration-options

Syntax **protocol-configuration-options** {**apco** | **pco**}
no protocol-configuration-options

Context config>subscr-mgmt>gtp>peer-profile

Description This command configures the Information Element to use for the Protocol Configuration Options.

The **no** form of this command reverts to the default value.

Default protocol-configuration-options pco

Parameters **apco** — Specifies that the system uses the Protocol Configuration Options Information Element.

pco — Specifies that the system uses the Additional Protocol Configuration Options Information Element.

python-policy

Syntax **python-policy** *name*
no python-policy

Context config>subscr-mgmt>gtp>peer-profile

Description This command specified the python-policy for MGW profile packets sent/received on the group interface

The **no** form of this command removes the policy name from the configuration.

Default no python-policy

Parameters *name* — Specifies an existing Python policy name up to 32 characters.

rat-type

Syntax **rat-type** {**utran** | **geran** | **wlan** | **gan** | **hspa** | **eutran** | **virtual** | *id*}
no rat-type

Context config>subscr-mgmt>gtp>peer-profile

Description This command configures the default Radio Access Type (RAT) signaled during GTP setup. RAT is the underlying physical connection method for a radio-based communication network. This can be overridden by RADIUS.

The **no** form of this command reverts to the default value.

Default rat-type wlan

Parameters **utran** — Specifies the signaled RAT type is UTRAN (1).

geran — Specifies the signaled RAT type is GERAN (2).

wlan — Specifies the signaled RAT type is WLAN (3).

gan — Specifies the signaled RAT type is GAN (4).

hspa — Specifies the signaled RAT type is HSPA Evolution (5).

eutran — Specifies the signaled RAT type is EUTRAN (6).

virtual — Specifies the signaled RAT type is virtual (7).

id — Specifies the numeric RAT type value.

Values 0 to 255

report-wlan-location

Syntax [**no**] **report-wlan-location**

Context config>subscr-mgmt>gtp>peer-profile

Description This command enables reporting the WLAN location or cellular location of the UE in the signaling interface (S2a or Gn) between the WLAN GW and the mobile gateway (PGW or GGSN).

The **no** form of this command disables location reporting.

Default no report-wlan-location

session-hold-time

Syntax **session-hold-time remaining-lease-time**
session-hold-time *seconds*
no session-hold-time

Context config>subscr-mgmt>gtp>peer-profile

Description This command configures, in seconds, the time that a GTP session context is held after the corresponding User Equipment is disconnected. If the same User Equipment re-connects to this system before this time has elapsed, its GTP session context is re-used. When the timer expires, the GTP session context is cleared.

The **no** form of this command reverts to the default.

Default session-hold-time 30

Parameters **remaining-lease-time** — Specifies that the timer is equal to the User Equipment's DHCP remaining lease time.

seconds — Specifies the time, in seconds, to hold a GTP session after its UE is disconnected.

Values 0 to 3600

serving-network

Syntax **serving-network mcc** *mcc-value* **mnc** *mnc-value*
no serving-network

Context config>subscr-mgmt>gtp

Description This command configures the Operator Identifier part (MCC and MNC) of the APN.

The **no** form of this command removes the values from the profile.

Default no serving-network

Parameters *mcc-value* — Specifies the Mobile Country Code (MCC) portion of the Serving Network.

Values 3 digits

mnc-value — Specifies the Mobile Network Code (MNC) portion of the Serving Network.

Values 2 or 3 digits

12.5.2.1.3 GTP Show Commands

gtp

Syntax **gtp**

Context show>router

Description This command displays operational information related to GTP usage in the specified routing context.

peer

Syntax **peer** [**peer-profile** *profile-name*] [**local-address** *ip-address*] [**control** *protocol*] [**interface-type** *interface-type*]
peer remote-address {*ip-address* | *ipv6-address*] [**udp-port** *port*] [**dsm-local-ip-address** {*ip-address* | *ipv6-address*}
peer remote-address {*ip-address* | *ipv6-address*] [**udp-port** *port*] **statistics** [**dsm-local-ip-address** {*ip-address* | *ipv6-address*}

Context show>router>gtp

Description This command displays GTP peer information.

Parameters *profile-name* — Specifies the name that identifies the profile.

ip-address — Specifies the local IP address.

Values

ip-address:	ipv4-address - a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D

protocol — Specifies the control plane protocol used for the connection with this Mobile Gateway.

Values gtpv1-c, gtpv2-c

interface-type — Specifies the interface type used in the context of this peer.

Values gn — Gn interface
s2a — S2a interface
s2b —S2b interface
s11 —S11 interface

ip-address — Specifies the remote IP address.

Values

ip-address:	ipv4-address - a.b.c.d
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D

port — Specifies the UDP port.

Values 1 to 65535

statistics — Displays statistics information about the Mobile Gateways connected to this system.

Output The following is an example of peer GTP information.

Sample Output

```
Node# show router 10 gtp peer
===============================================================================
Peers
===============================================================================
Remote address              : 192.168.100.1
UDP port                    : 2123
-------------------------------------------------------------------------------
State                       : up
Local address               : 192.0.2.10
Profile                     : default_s11
Control protocol            : gtpv2-c
Interface type              : s11
Restart count               : 1
Time                        : 2018/05/31 11:12:59
-------------------------------------------------------------------------------
No. of Peers: 1
===============================================================================
```

peer-profile-map

Syntax **peer-profile-map**

Context show>router>gtp>s11

Description This command displays the peer profile mapping that is used to connect to S11 peers.

Output The following is an example of GTP peer profile map information.

Sample Output

```
Node# show router 20 gtp s11 peer-profile-map
===============================================================================
Peer profile map
===============================================================================
Address prefix                              Profile
-------------------------------------------------------------------------------
0.0.0.0/0                                   default_s11
192.2.0.0/24                                specific_profile
-------------------------------------------------------------------------------
No. of address prefixes: 2
===============================================================================
```

uplink

Syntax **uplink**
uplink mgw-address-cache [**arec**] [**snaptr**] [**srv**]
uplink mgw-address-cache apn *apn-domain-string*

Context show>router>gtp

Description This command displays operational information related to the usage of GTP uplink interfaces (Gn, S2a, S2b) in this routing context.

Parameters **arec** — Displays A-records.

snaptr — Displays Straightforward-NAPTR information.

srv — Displays SRV records.

apn-domain-string — Specifies the Access Point Name (APN) of this DNS cache entry.

Output The following is an example of GTP uplink information.

Sample Output

```
Node# show router 10 gtp uplink mgw-address-cache
===============================================================================
Mobile Gateway address cache
===============================================================================
APN : full.dotted.apn.apn.epc.mnc010.mcc206.3gppnetwork.org
-------------------------------------------------------------------------------
```

```
Mobile Gateway address : 5.20.1.2
Time left (s) : 3587
-------------------------------------------------------------------------------
No. of cache entries: 1
No. of Mobile gateways: 1
===============================================================================
```

gtp

Syntax **gtp**

Context show>subscr-mgmt

Description This command displays router operational GTP information.

peer-profile

Syntax **peer-profile** [*profile-name*]
peer-profile *profile-name* **associations**

Context show>subscr-mgmt>gtp

Description This command displays GTP peer profile information.

Parameters *profile-name* — Specifies the name that identifies the profile.

associations — Specifies the associated contexts for the specified peer profile.

Output The following is an example of GTP peer profile information.

Sample Output

```
Node# show subscriber-mgmt gtp peer-profile
===============================================================================
GTP peer profiles
===============================================================================
Name                            Description
-------------------------------------------------------------------------------
default                          (Not Specified)
default_s11                      (Not Specified)
-------------------------------------------------------------------------------
No. of profiles: 2
===============================================================================
Node# show subscriber-mgmt gtp peer-profile "default_s11"
===============================================================================
GTP peer profile "default_s11"
===============================================================================
Description                 : (Not Specified)
Retransmit timeout (s)      : 5
Retransmit retries          : 3
Keepalive interval (s)      : 60
Keepalive retries           : 4
```

```
Keepalive retry timeout (s) : 5
Time to live                : 255
IPv4 MTU                    : 1400
Interface type              : s11
Charging char home UE       : (None)
Charging char roaming UE    : (None)
Session hold time (s)       : 30
Report WLAN location        : disabled
Procotol Config Options I.E.: pco
Python policy               : (Not Specified)
Radio Access Technology type: wlan
Last management change      : 05/31/2018 11:05:46
GGSN uplink GBR (Kbps)      : 5000
GGSN uplink MBR (Kbps)      : 5000
GGSN downlink GBR (Kbps)    : 2000
GGSN downlink MBR (Kbps)    : 2000
GGSN Alloc/Retention Prio   : 1
GGSN uplink AMBR (Kbps)     : (Not Specified)
GGSN downlink AMBR (Kbps)   : (Not Specified)
GGSN last management change : 05/31/2018 11:05:46
PGW uplink GBR (Kbps)       : 0
PGW uplink MBR (Kbps)       : 0
PGW downlink GBR (Kbps)     : 0
PGW downlink MBR (Kbps)     : 0
PGW Alloc/Retention Prio    : 1
PGW Qos Class ID            : 8
PGW uplink AMBR (Kbps)      : 10000
PGW downlink AMBR (Kbps)    : 20000
PGW last management change  : 05/31/2018 11:05:46
MME uplink GBR (Kbps)       : 0
MME uplink MBR (Kbps)       : 0
MME downlink GBR (Kbps)     : 0
MME downlink MBR (Kbps)     : 0
MME Alloc/Retention Prio    : 1
MME Qos Class ID            : 8
MME uplink AMBR (Kbps)      : 10000
MME downlink AMBR (Kbps)    : 20000
MME last management change  : 05/31/2018 11:05:46
===============================================================================
```

s11

Syntax **s11**

s11 session imsi *imsi* **apn** *apn-string*

s11 session [**peer-address** *ip-address*] [**router** *router-instance*] [**remote-control-teid** *teid*] [**local-control-teid** *teid*] [**detail**]

s11 session imsi *imsi*

Context show>subscr-mgmt>gtp

Description This command displays information related to established S11 connections. Specific sessions can be selected based on IMSI, APN, or MME.

Parameters *imsi* — Specifies the IMSI (International Mobile Subscriber Identity) of this UE.

apn-string — Specifies the APN (Access Point Name).

ip-address — Specifies the IP address of the MME.

router-instance — Specifies the identifier of the virtual router instance where the GTP tunnel is terminated.

remote-control-teid *teid* — Specifies the remote control plane Tunnel Endpoint Identifier (TEID).

local-control-teid *teid* — Specifies the local control plane TEID.

detail — Displays detailed information.

Output The following is an example of GTP S11 information.

Sample Output

```
Node# show subscriber-mgmt gtp s11 session
===============================================================================
GTP S11 sessions
===============================================================================
IMSI                        : 001020000000001
APN                         : demo.mnc001.mcc001.gprs
-------------------------------------------------------------------------------
Peer router                 : 10
Peer address                : 192.168.100.1
Remote control TEID         : 1
Local control TEID          : 4293919008
PDN TEID                    : 4293919008
Charging characteristics    : (None)
Uplink AMBR (kbps)          : 10000
Downlink AMBR (kbps)        : 20000
Ipoe-session SAP            : [pxc-1.b:1.8]
Ipoe-session Mac Address    : 00:03:ff:f0:01:20
-------------------------------------------------------------------------------
No. of GTP S11 sessions: 1
===============================================================================
```

statistics

Syntax **statistics**

Context show>subscr-mgmt>gtp

Description This command displays GTP statistics information.

Output The following is an example of GTP statistic information.

Sample Output

```
Node# show subscriber-mgmt gtp statistics
===============================================================================
GTP statistics
===============================================================================
```

```
tx echo requests                                : 23
tx echo responses                               : 0
tx errors                                       : 0
rx echo requests                                : 0
rx echo responses                               : 23
rx errors                                       : 0
rx version not supported                        : 0
rx zero TEID responses                          : 0
path faults                                     : 0
path restarts                                   : 0
tx invalid msgs                                 : 0
tx create PDP context requests                  : 0
tx create PDP context responses                 : 0
tx delete PDP context requests                  : 0
tx delete PDP context responses                 : 0
tx create session requests                      : 0
tx create session responses                     : 1
tx delete session requests                      : 0
tx delete session responses                     : 0
tx delete bearer requests                       : 0
tx delete bearer responses                      : 0
tx create bearer responses                      : 0
tx update bearer responses                      : 0
tx modify bearer requests                       : 0
tx modify bearer responses                      : 1
tx change notification responses                : 0
tx rls access bearers responses                 : 0
tx error indication count                       : 0
tx downlink notifications                       : 0
tx stop paging indication                       : 0
tx end marker                                   : 0
rx invalid msgs                                 : 0
rx create PDP context requests                  : 0
rx create PDP context responses                 : 0
rx delete PDP context requests                  : 0
rx delete PDP context responses                 : 0
rx create session requests                      : 1
rx create session responses                     : 0
rx delete session requests                      : 0
rx delete session responses                     : 0
rx delete bearer requests                       : 0
rx delete bearer responses                      : 0
rx create bearer requests                       : 0
rx update bearer requests                       : 0
rx modify bearer requests                       : 1
rx modify bearer responses                      : 0
rx change notification requests                 : 0
rx rls access bearers request                   : 0
rx error indication count                       : 0
rx downlink notification ack                    : 0
rx downlink notification fail                   : 0
rx invalid pkt length                           : 0
rx unknown pkts                                 : 0
rx missing IE pkts                              : 0
rx bad UDP header pkts                          : 0
rx discarded pkts                               : 0
rx in-session discarded pkts                    : 0
rx pkts                                         : 25
tx discarded pkts                               : 0
```

```
tx pkts                                                 : 25
rx UDP checksum error                                   : 0
===============================================================================
```

summary

Syntax **summary**

Context show>subscr-mgmt>gtp

Description This command displays a GTP system summary of information.

Output The following is an example of GTP summary information.

Sample Output

```
Node# show subscriber-mgmt gtp summary
===============================================================================
GTP system summary
===============================================================================
Actual number of MME                                    : 1
Actual number of ENODE-B                                : 1
Actual number of S11 Sessions                           : 1
Actual number of Mobile Gateways                        : 0
Actual number of Uplinks                                : 0
Actual number of Uplinks in Hold                        : 0
===============================================================================
```

uplink

Syntax **uplink**
uplink session imsi *imsi* **apn** *apn-string*
uplink session [**peer-address** *ip-address*] [**router** *router-instance*] [**remote-control-teid** *teid*] [**local-control-teid** *teid*] [**detail**]
uplink session imsi *imsi*

Context show>subscr-mgmt>gtp

Description This command displays box-wide operational information related to the GTP uplink interfaces (Gn, S2a, S2b).

Parameters *imsi* — Specifies the IMSI (International Mobile Subscriber Identity) of this UE.

apn-string — Specifies the APN (Access Point Name).

ip-address — Specifies the IP address of the Mobile Gateway (PGW or GGSN).

router-instance — Specifies the identifier of the virtual router instance where the GTP tunnel is terminated.

remote-control-teid *teid* — Specifies the remote control plane Tunnel Endpoint Identifier (TEID).

local-control-teid *teid* — Specifies the local control plane TEID.

detail — Displays detailed information.

Output The following is an example of GTP uplink information

Sample Output

```
Node# show subscriber-mgmt gtp uplink session detail
===============================================================================
GTP sessions
===============================================================================
IMSI : 206100000000041
APN : full.dotted.apn.mnc010.mcc206.gprs
-------------------------------------------------------------------------------
Mobile Gateway router : "Base"
Mobile Gateway address : 5.20.1.2
Remote control TEID : 1119232
Local control TEID : 4293918976
Bearer 5 rem TEID : 1074861061
Bearer 5 loc TEID : 4293919013
-------------------------------------------------------------------------------
No. of GTP sessions: 1
===============================================================================
```

12.5.2.1.4 GTP Clear Commands

gtp

Syntax **gtp**

Context clear>router

Description This command enables the context to clear GTP information in this routing context.

peer

Syntax **peer remote-address** *ip-address* [**remote-udp-port** *port*]
peer remote-address *ip-address* [**remote-udp-port** *port*] **statistics**

Context clear>router>gtp

Description This command clears all states for a given peer in this routing context. The system no longer sends keep-alive messages to the peer, and all active sessions are terminated.

If the optional statistics parameter is specified, only the peer statistics are cleared and all sessions remain.

Parameters *ip-address* — Specifies the IP address of the peer.

Values

ipv4-address:	a.b.c.d (host bits must be 0)
ipv6-address:	x:x:x:x:x:x:x:x (eight 16-bit pieces)
	x:x:x:x:x:x:d.d.d.d
	x - [0 to FFFF]H
	d - [0 to 255]D

port — Specifies the UDP port used by the peer.

statistics — Specifies to clear only the peer statistics with all sessions remaining active.

uplink

Syntax **uplink**

Context clear>router>gtp

Description This command clears information related to the GTP uplink state for the routing context.

mgw-address-cache

Syntax **mgw-address-cache**

Context clear>router>gtp>uplink

Description This command clears the DNS resolution cache for all APNs known in this router. After execution, every APN to address mapping is learned again using DNS.

gtp

Syntax **gtp**

Context clear>subscr-mgmt

Description This command clears box-wide GTP state information.

session

Syntax **session all**

session apn *apn*
session apn *apn* **imsi** *imsi*
session imsi *imsi*

Context clear>subscr-mgmt>gtp

Description This command clears all GTP sessions matching the specified criteria. Associated objects, such as the ESM subscriber in a GTP uplink scenario are also removed.

Parameters *apn* — Specifies the APN identifying the session.

imsi — Specifies the IMSI identifying the session.

all — Specifies to remove all sessions in the system.

statistics

Syntax **statistics**

Context clear>subscr-mgmt>gtp

Description This command clears all box-wide GTP statistics.

12.5.2.1.5 GTP Debug Commands

gtp

Syntax **gtp**

Context debug

Description This command enables the context to configure debugging for GTP.

event

Syntax [**no**] **event**

Context debug>gtp

Description This command configures detailed debugging of all events in the GTP system.

The **no** form of this command disables event debugging.

imsi

Syntax [**no**] **imsi** *imsi*

Context debug>gtp

Description This command restricts debugging to only data related to the specified IMSI. This command can be repeated multiple times, where only data for any of the specified IMSIs is debugged.

The **no** form of this command removes the filter for the specified IMSI. If the last IMSI filter is removed, all data is debugged again, but may be restricted by other filters.

Parameters *imsi* — Specifies the mobile subscriber identity, as a string of up to 15 characters.

packet

Syntax **packet detail-level** {**high** | **low**}
packet mode {**all** | **dropped**}
no packet

Context debug>gtp

Description This command enables debugging of GTP packets sent or received by the system's control plane.

The **no** form of this command disables debugging of GTP packets.

Parameters **detail-level** {**high** | **low**} — Specifies how much detail is to be displayed when debugging a GTP packet.

Values

high	Specifies to display and decode all data in the packet.
low	Specifies to display and decode only the most important data.

Default low

mode {**all** | **dropped**} — Specifies which packets is debugged.

Values

all	Specifies to debug all packets.
dropped	Specifies to debug only dropped packets.

Default dropped

peer

Syntax [**no**] **peer router** *router-instance* **address** *ip-address* [**udp-port** *port*]

Context debug>gtp

Description This command restricts debugging to only data related to the specified GTP peer. This command can be repeated multiple times, where only data for any of the specified peers is debugged.

The **no** form of this command removes the restriction for the specified peer. When the last peer filter is removed, all data is debugged again, but may be restricted by other filters.

Parameters *router-instance* — Specifies the ID of the VRF where the peer is connected.

Values

router-instance:	*router-name* | *vprn-svc-id*
router-name:	"Base"
vprn-svc-id	1 to 2147483647

ip-address — Specifies the IP address of the peer.

Values a.b.c.d

port — Specifies the GTP-C port used by the peer.

13 Virtual Residential Gateway

This section describes virtualized residential gateway (vRGW).

13.1 Overview

A vRGW transforms the Layer 3 routed Residential Gateway (RGW) in the home into a Bridged Residential Gateway (BRG), by moving the Layer 3 functions of an RGW into the network resides on a node (vRGW). The Layer 3 functionalities (such as address management, routing, Internet connectivity, NAT, UPnP, firewall, and application awareness) provides functions such as URL filtering and parental control and is moved to the network and resides on the vRGW. The RGW then operates in a bridged mode and performs local switching of intra-home traffic that originates and terminates on devices within the home. Bridged traffic destined for outside the home (such as to the Internet, the provider's content, or another home) over the WAN toward the vRGW is hereby referred to as a BRG.

This mode of operation allows an operator visibility of the connected devices on the home LAN, instead of just a single IP address per home as is the case for a Layer 3 RGW. This allows operational improvements through per-device control and troubleshooting, and the ability to offer new services faster and on a more granular device specific-basis.

The router, as a dedicated gateway or as a BNG, serves as a vRGW providing Layer 3 termination and ESM functionality for bridged homes, as shown in Figure 193.

Figure 193 **Virtualized Residential Gateway (vRGW)**

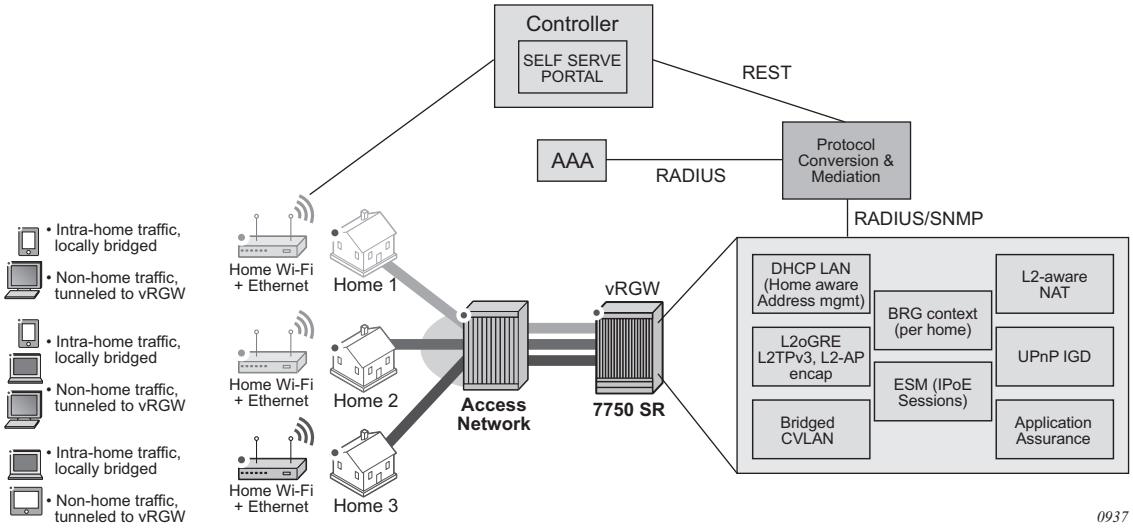

13.1.1 Access Modes

All vRGW functionality is supported on both regular group interfaces (SAPs, LAGs, PW-SAPs, and so on) and WLAN GW group interfaces (soft-GRE, soft-L2TPv3, VLAN).

A new configuration sub-node for the BRG is provided under the **group-interface** context for regular group interfaces and under the **vlan-range** context for WLAN-GW group interfaces. A regular group interface with a BRG sub-node does not support any non-BRG configuration and must operate in the **ipoe-bridged** mode.

In a WLAN-GW group interface, the BRG is configured in the **vlan-range** level. With a **vlan-range** it is not possible to mix BRG and other existing functionalities, but it is possible to mix BRG and other functionalities (such as WLAN-GW) on the same group interface. If a BRG also supports public WIFI, the expectation is that the BRG will have different SSIDs for public WIFI and for private home traffic on the BRG, each represented by a different VLAN tag.

Contrary to WLAN-GW UEs, which require anchoring based on their MAC addresses (for mobility), devices associated with a BRG are anchored based on the tunnel source IP address of the BRG. The system therefore load-balances on a per-BRG granularity basis across a set of configured ISAs. Anchoring based on the source IP address of the BRG allows all devices in the home to be anchored on the same ISA and IOM. This enables aggregate QoS functionality within a single home.

Tunnel QoS is not supported as this is performed by regular subscriber QoS in the BRG scenario. WLAN-GW IOM (N:M) redundancy is supported. Data-triggered authentication (IPv4 only) is supported. All WLAN-GW access types (GRE, L2TPv3, L2-AP) are supported.

13.1.2 Home Context on the vRGW

The system keeps a management context for each connected home or BRG, which is identified by a BRG ID. This context is used for authentication, configuration, and retrieval of operational data. Authentication can be performed explicitly by the BRG, in which case the vRGW acts as a RADIUS proxy. Alternatively, the vRGW can implicitly authenticate on the BRG's behalf when the first host from the BRG connects.

Initial configuration parameters are provided as RADIUS attributes during authentication. Configuration parameters provided on the home level is used as defaults on the host level can be overridden on a per-host basis. Home-level configuration can be dynamically overridden by means of a RADIUS CoA message. Refer to the "Virtual Residential Gateway" section in the *7750 SR and VSR RADIUS Attributes Reference Guide*.

Operational data provides information about the BRG, such as the number of connected devices, how long the BRG has been active, and which configuration parameters were provided. This can be retrieved via MIBs or show commands.

The vRGW communicates with a classic AAA server mainly to perform the authentication and with a separate configuration or operational platform. This separate configuration/operational platform is referred to as a "controller" and can be combined in a single management control system.

13.1.2.1 Implicit Home Authentication

With implicit home authentication (see Figure 194), the vRGW authenticates a new BRG when the first associated device connects. To avoid restrictions on the connectivity model, the vRGW initially attempts to identify a BRG with a BRG ID. The vRGW always performs authentication of a new host, and this authentication should return the BRG ID with which this host is associated. If this BRG ID is not yet known, the vRGW triggers BRG-level authentication. This allows an operator flexibility to identify a home. For example, one deployment might use CVLANs as an identifier while another might use a BRG MAC as an identifier.

When using the BRG MAC, this can be learned using the same methods as the AP MAC for WLAN-GW, as described in WIFI Aggregation and Offload, Authentication. The learned MAC can subsequently be reflected in authentication of the new device and mapped to a suitable BRG-ID by the AAA server or controller.

While a home can be optionally authenticated with an AAA server, each device in the home, is typically only authorized with a controller (via regular RADIUS messages and configured authentication policy) to get its configuration. This authorization is handled by the controller, which identifies and returns the associated BRG ID and, optionally, any device-level configuration.

Per-home authentication is forwarded to the AAA server to be fully authenticated. In this case, the controller typically performs an AAA proxy functionality so it can insert home configuration data in the final Access-Accept message. After both the device and BRG authentication are completed, the resulting RADIUS attributes are used to set up all required ESM objects (hosts, subscribers, SLA profile instances, filters, and so on).

Figure 194 **Implicit Home Authentication**

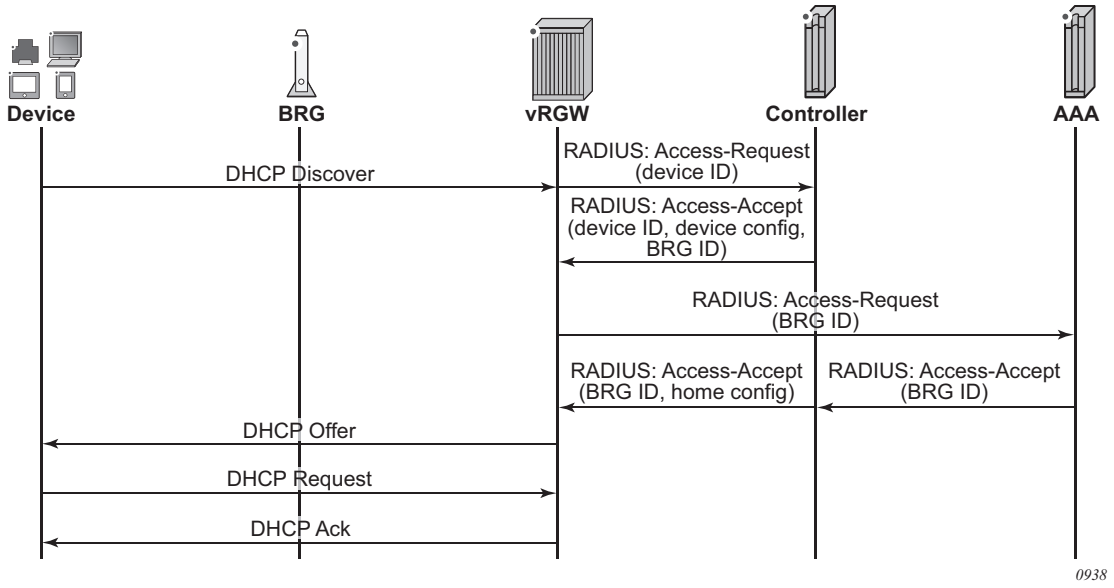

13.1.2.2 Explicit Home Authentication

With explicit home authentication (see Figure 195), the BRG authenticates itself at boot time before allowing connectivity of any devices. The vRGW acts as a proxy for this authentication, and upon receipt of a successful authentication, the BRG context is created and stores all home parameters. As with implicit authentication, the vRGW performs authentication for each device and expects the BRG ID as a result.

Figure 195 **Explicit Home Authentication**

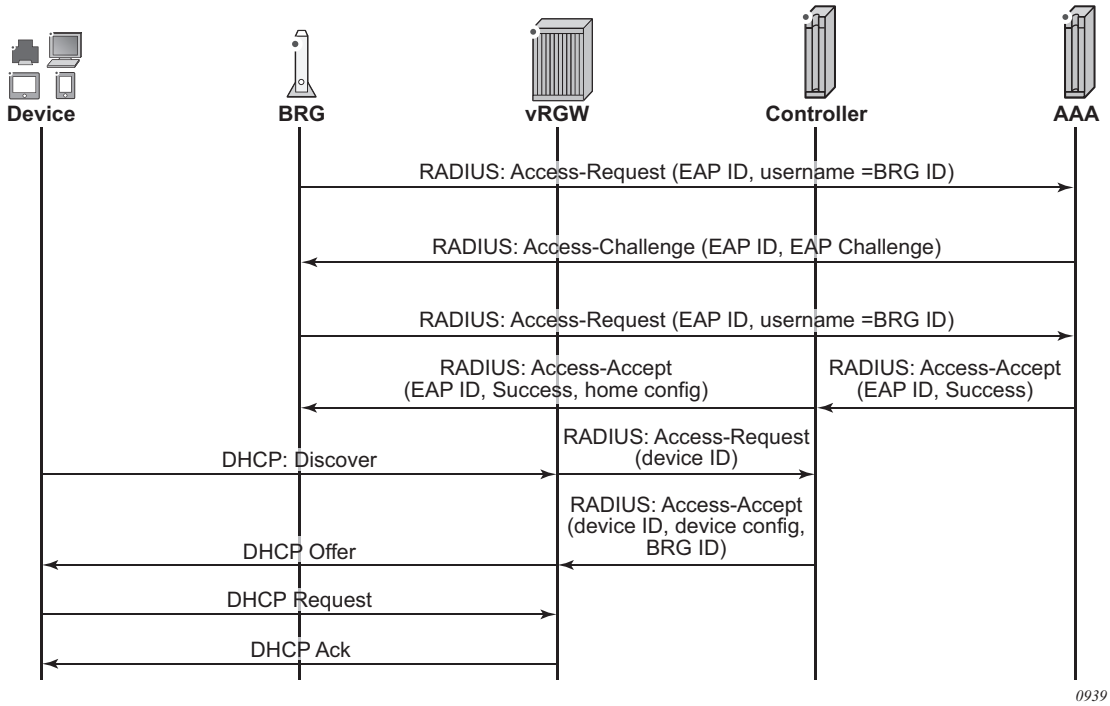

On the router, this functionality is obtained by including a RADIUS proxy in the vRGW configuration. This RADIUS proxy can be used for both WLAN-GW and vRGW scenarios and determines the difference based on the presence of a BRG ID attribute in the Access-Accept message. Proxy transactions in the context of a BRG are not cached. Instead, a fully functional BRG context is already created. Track accounting is not supported in the context of a BRG. Re-authentication is supported and behaves the same as on a CoA at the BRG level.

13.1.2.3 Change of Configuration

RADIUS CoA can be used to change the configuration on a home level. The CoA can use the BRG ID as a key and contain attributes for all new or updated parameters. For all correctly specified parameters, the vRGW will overwrite the existing configuration on the home level and updates all devices connected to the BRG.

13.1.2.4 Home Lifetime

A configurable ping based on a smart connectivity-verification mechanism is provided to detect whether a BRG is alive. The BRG state is removed if it is not deemed alive (subject to a configurable hold time). A BRG is always considered to be alive as long as there are connected non-static devices. Static devices are not considered because there is no control plane to track whether or not they are alive. When no sessions are present, the home context can be in one of the following conditions.

- On explicit authentication, where there is no session connection, the BRG state should be kept alive. To control this, an initial hold time applies during which the BRG context cannot be removed. Any connected device cancels the timer. If the timer expires and there are still no connections present, the BRG fall backs to the behavior as if the last session was removed.
- If the last session was removed and connectivity verification is configured, the vRGW tries to ping the BRG. This ping mechanism uses ICMP or ICMPv6 ping toward the tunnel source IP or, if not present, toward the RADIUS source IP address. If the pinging is successful, the BRG context is kept. If the pinging is unsuccessful or impossible to perform (for example, if no IP address is known), the BRG falls back to the hold time. If, during the connectivity verification, a host connects to the BRG, the verification is stopped and the BRG is assumed to be alive again.
- If the last session was removed and connectivity verification failed or was not configured, a configurable hold time applies, during which time the BRG is not removed. If a new device connects during a hold time, the timer is canceled. If the timer is not canceled, the BRG is removed when the timer expires. The configured hold time can be zero and can be different from the initial hold time.

13.1.3 Device Context on the vRGW

The router maps every device in a home to an IPoE session, which can contain one or more ESM hosts, depending on the number of IP stacks active on the device. The vRGW authorizes every single device and stores the resulting configuration data with the IPoE session. Whenever the session is created or updated, the vRGW combines the configuration data of the home context with the configuration data of the device. When the same configuration object is present on both levels, the device level applies. The resulting combined configuration data is used to install or update the ESM objects.

13.1.4 Dynamic Configuration Changes

Use a RADIUS CoA to dynamically change the configuration on both the home level and the device level. Use the BRG ID as the key on the CoA on the home level. The included attributes overwrites the existing stored configuration on the home level and subsequently update all devices connected to the home. All devices pick up the new home-level configuration unless a more specific configuration exists on the device level.

A CoA on the device level can use all existing ESM keys as detailed in the 7750 SR and VSR RADIUS Attributes Reference Guide. The keys update the existing configuration on the device level. If the device inherits the specified parameter from the home level, this parameter is updated on the device level; the existing home-level configuration remains intact and is used as the default for other devices.

In certain cases, the configuration can be removed on the device level to be able to fall back to the home level, supported by the RADIUS attribute Alc-Remove-Override. This attribute lists which overrides must be removed on the session level and fall back (revert) to the home-level configuration. If the home parameter changes, the device also picks up the new configuration.

For more details on RADIUS configuration attributes, refer to the *7750 SR and VSR RADIUS Attributes Reference Guide*.

13.1.5 Per-Home Pool Management and L2-Aware NAT

This feature allows the provisioning of a DHCP pool per home in a vRGW context. The addresses in a per-home pool are unique, so local bridging can occur within the home functions. Devices in different homes can, however, be allocated the same address. L2-Aware NAT is then used to handle address translation and connectivity toward the network and between homes. Per-home default GW, subnet, and address range (out of which addresses are allocated to home devices) can be configured via the CLI. The address range can be overridden from RADIUS. The subnet associated with the home pool must be within the configured L2-Aware inside subnet. L2-Aware source NAT can optionally be combined with destination NAT to support enhanced traffic redirection (such as for stateful DNS overwrite function). More details on network address translation forwarding can be found in the *7450 ESS, 7750 SR, and VSR Multiservice Integrated Service Adapter Guide*.

In addition to the IPv4 addresses from the home pool, the following hosts can be set up:

- IPv4 DHCP proxy hosts using an AAA-provisioned address

  Only non L2-Aware hosts can receive a framed IP address from RADIUS.
- IPv6 SLAAC hosts using an AAA-provisioned address or a local address assignment
- IPv6 IA_NA hosts using DHCPv6 relay, an AAA-provisioned address, or a local address assignment

13.1.5.1 Sticky IP Addresses

The vRGW can be programmed with a list of sticky IP addresses per BRG. A sticky IP address is an IPv4 address from the home pool allocation range that is set aside for a specific device based on the MAC address and is never allocated to any other device. This can be used for devices where a persistent IP address is desirable but configuring a static IP address on the device is too cumbersome (such as a network printer). The device will still use DHCP to gain connectivity but is always assigned the same IP address.

Public (non-NAT) sticky IP addresses are not directly supported in pool management but can be provisioned in the following manner.

1. Create a local DHCP server and pool to manage the public address space.

2. For each public sticky IP address that must be provisioned, create a sticky lease in this pool. Refer to DHCP Address Reservation for Sticky Leases for further details. Both the user identification and hostname can be the MAC address of the associated device.
3. When the device linked to the sticky lease authorizes with the controller, it returns the allocated address as a framed IP address. This address takes precedence over any home pool configuration and a DHCP proxy host is installed.

13.1.5.2 Managed Static IPv4 Addresses

Similar to sticky IP addresses, the vRGW can be programmed with a list of static IP addresses per BRG. However, for static IP addresses, DHCP is not expected from clients and is dropped. The vRGW will automatically install all of these addresses as IPv4 hosts contained in IPoE sessions. Because every host must be linked to a specific point of access (such as a SAP or tunnel) the vRGW needs to wait on a data trigger or the first non-static host in a home before static hosts can be created. Similar to regular (DHCP-triggered) hosts, a device-level authorization is performed for each static host to retrieve the per-device level configuration.

This type of static host can cover both private (inside the home subnet) and public (non-L2-Aware NAT) addresses. The created hosts can only be explicitly removed by the management interface, and mechanisms such as session-timeout and idle-timeout are not supported.

13.1.5.3 DMZ

A single DMZ address can be provisioned per home via RADIUS (VSA Alc-DMZ-Address), which can be any address inside the home subnet except for the default gateway. When a host exists with this address, the DMZ mode is activated. DMZ is be activated if a host exists with the address and the subscriber uses only one port range per IP. Without DMZ mode enabled, any traffic arriving for the NAT outside IP that does not match an existing flow, pinhole, or port-forward, is dropped. With DMZ mode enabled, this traffic is forwarded to the provisioned DMZ host.

13.1.6 IPv6

The vRGW supports the following IPv6 host types:

- IPv6 SLAAC hosts using an AAA-provisioned address or a local address assignment
- IPv6 IA_NA hosts using DHCPv6 relay, an AAA-provisioned address, or a local address assignment

The vRGW only operates in IPoE bridge mode. For regular group interfaces, the IPoE bridge mode must be explicitly enabled before the BRG can be enabled. For WLAN-GW group interfaces, the IPoE bridge mode is implicitly assumed for a VLAN with BRG enabled. This has the following consequences:

- SLAAC hosts from the same home can share a /64 prefix
- when a local address assignment is used, SLAAC hosts from the same home are automatically assigned the same /64 prefix
- IA_NA hosts from the same BRG can receive unique addresses from a shared /64 prefix. This prefix is automatically signaled in an RA

13.1.7 QoS and Filter Support

The vRGW is based on existing ESM QoS configurations on the router where each home maps to a single subscriber instance. Home-level bandwidth can be provided by an aggregate rate or scheduler policy on the subscriber level. Groups of home devices (such as telephones, computers, televisions, and security systems) can be given a shared bandwidth by assigning a different SLA profile per such group. Individual device-level QoS can be provisioned by mapping a specific device to a specific forwarding class using IP classifiers based on the device IP. This forwarding class can then be mapped to its own individual queue.

Dynamic QoS overrides can be provisioned on both the home level and the device level. While QoS overrides are represented by a single RADIUS attribute, the device level does not override the whole attribute but only the QoS objects specified on that level. For example, on the home level, the scheduler bandwidth is overridden, and on the device level, queue bandwidth is overridden. The resulting override is a combination of both.

13.1.8 Data-Triggered Authentication

ISA-based IPv4 data-triggered authentication (Figure 196) and host creation is supported on WLAN-GW group interfaces. When authenticating a data-triggered device, connection data provides less data for the controller to derive the BRG, unlike DHCP where the BRG can insert its identifier, such as a circuit ID. For pure Layer 2 access, a BRG ID can be hard-coded to a port and VLANs, although for tunneled access, this is not always possible as the corresponding value would be the tunnel source IP address. This IP address can be dynamically assigned and changed with a BRG reboot. The following alternatives are suggested.

- Use AP MAC as the identifier. This identifier can be signaled in DHCP and DHCPv6 as specified in WIFI Aggregation and Offload. For data-trigger purposes, it can also be sent as part of the L2TPv3 header, or if a GRE, it can be learned upon data-trigger via an ARPoGRE/NDoGRE message.
- Use a custom identifier that is sent in DHCP in a circuit ID option or a vendor-specific option. To handle the data trigger while a DHCP lease is active, a controller keeps its state to map the device MAC to the BRG identifier.

If the data-triggered device is the first device to come up in the home and the BRG did not perform explicit authentication, the vRGW also triggers an implicit authentication. After authentication, the data-triggered host can be installed by one of the following methods.

- If it the data trigger was for a static IP address (for example, when the static device is the first to send upstream data in the home), the triggering static host and any other provisioned static hosts are installed.
- If the trigger was sent for a dynamic host (sticky/not sticky), (for example, when connection with a device was lost (based on an idle-timeout) but the lease was still valid), a DHCP lease is re-created using the provisioned lease time on the home level. This installed lease time is usually excessive compared to the actual remaining lease time on the device, but this is corrected when the lease performs DHCP renew or rebind procedures.

  The actual remaining lease time is used if known. If a host goes idle and sends a data trigger, the actual remaining lease time is used.

Figure 196 **Data-Triggered Authentication**

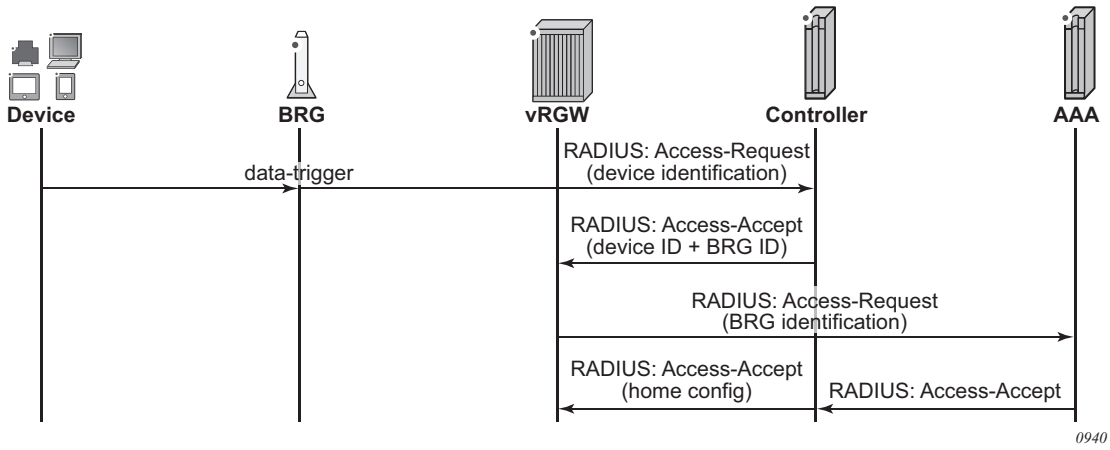

13.1.9 Per-Host NAT Port Ranges

Carriers want to offer opt-in value-added services (VASs) through dedicated DPI-based appliances or VMs in data centers. This functionality requires vRGW support to forward traffic (ideally, only for subscribers who have subscribed to the VAS) to the external appliance. The appliance implements per-subscriber and per-device policies, and must be able to determine the subscriber and device from the received packets. Because address space across homes can overlap, subscriber-aware NAT is a requirement in the vRGW architecture. When subscriber-aware NAT is used, the outside IP address is unique and corresponds to the subscriber but, by default, the device information is lost. However, the device can be determined for the external appliance from the Layer 3 packet if a unique NAT outside port range is used per device on the vRGW.

By default, the subscriber-aware NAT allows the entire port range (other than the port range for static port forwarding) to be available for dynamic NAT flows and dynamic port forwarding (via UPnP). The feature adds support for allocating per-host NAT outside port ranges, and reporting per-host port-range allocation and deallocation in RADIUS accounting. External VAS appliances can then track RADIUS accounting to determine device to port-range mapping.

The port range for a host is allocated and deallocated when the host is created and deleted, respectively. A single port range per host is supported. The RADIUS attribute **Alc-Per-Host-Port-Range** provides the count of ports per host for a subscriber. Refer to the *7750 SR and VSR RADIUS Attributes Reference Guide* for information about the attribute format. In case of multiple NAT policies per subscriber, the attribute value is required to be the same for all policies.

The presence of the VSA implicitly enables the per-host NAT port-range allocation mode. The **ports-per-host** mode is only enabled (via the VSA) if vRGW is enabled (as indicated by the presence of BRG in **no-shutdown**) under the VLAN range on the WLAN-GW interface, or on the group interface. The VSA can be present in **access-accept** for BRG authentication (implicit or explicit), and in CoA (with **Alc-BRG-Id** as the key). If a CoA is received with the **Alc-Per-Host-Port-Range** set to 0, it indicates the disabling of the per-host port-range mode.

If the **Alc-Per-Host-Port-Range** VSA is changed, the flows in the overlapping region between the new and old port ranges remain intact; remaining flows (if any) are removed.

- When the **ports-per-host** mode is disabled (via the VSA), the new port range includes all ports except the SPFs, and in accordance with the preceding rule, no ports are deleted.
- When **ports-per-host** mode is enabled mid-session (via the VSA), the new port range falls within the old port range, where the old range contained all ports except the SPFs.

The ISA is updated from the CPM for all hosts when the **ports-per-host** mode changes, and cleanup occurs in accordance with the preceding rule.

The **per-host port-range** is included in the **Alc-Nat-Port-Range** attribute in per-host or per-session RADIUS accounting (in accordance with the currently supported format of the **Alc-Nat-Port-Range** VSA for L2-Aware NATs).

When a new port block for a host is allocated or freed, an interim-update message with a **Nat-Port-Range-Event** reason is sent. The interim update is sent only when interim updates are enabled and the configured **include-attributes** contain the **alc-port-range** VSA.

The port range for a host remains allocated for the lifetime of the host (unless explicitly removed using the VSA). Per-host reserved ports (for prioritized sessions) and watermarks to indicate exhaustion of **per-host port-range** are not supported.

13.1.10 Inter-Chassis Redundancy

vRGW supports stateless inter-chassis redundancy, where the home and device state is not explicitly synchronized between the two vRGWs. After a redundancy switchover event, Data Triggered Authentication (DTA) is used to download the correct home and device state back from the controller.

The controller should always have the latest configuration state of a given home, including any overrides after the home was activated. Additionally, the controller maintains a minimal state for DHCP pool management that can be reflected to the vRGW in case of a redundancy event, as shown in Figure 197.

Figure 197 **Inter-Chassis Redundancy**

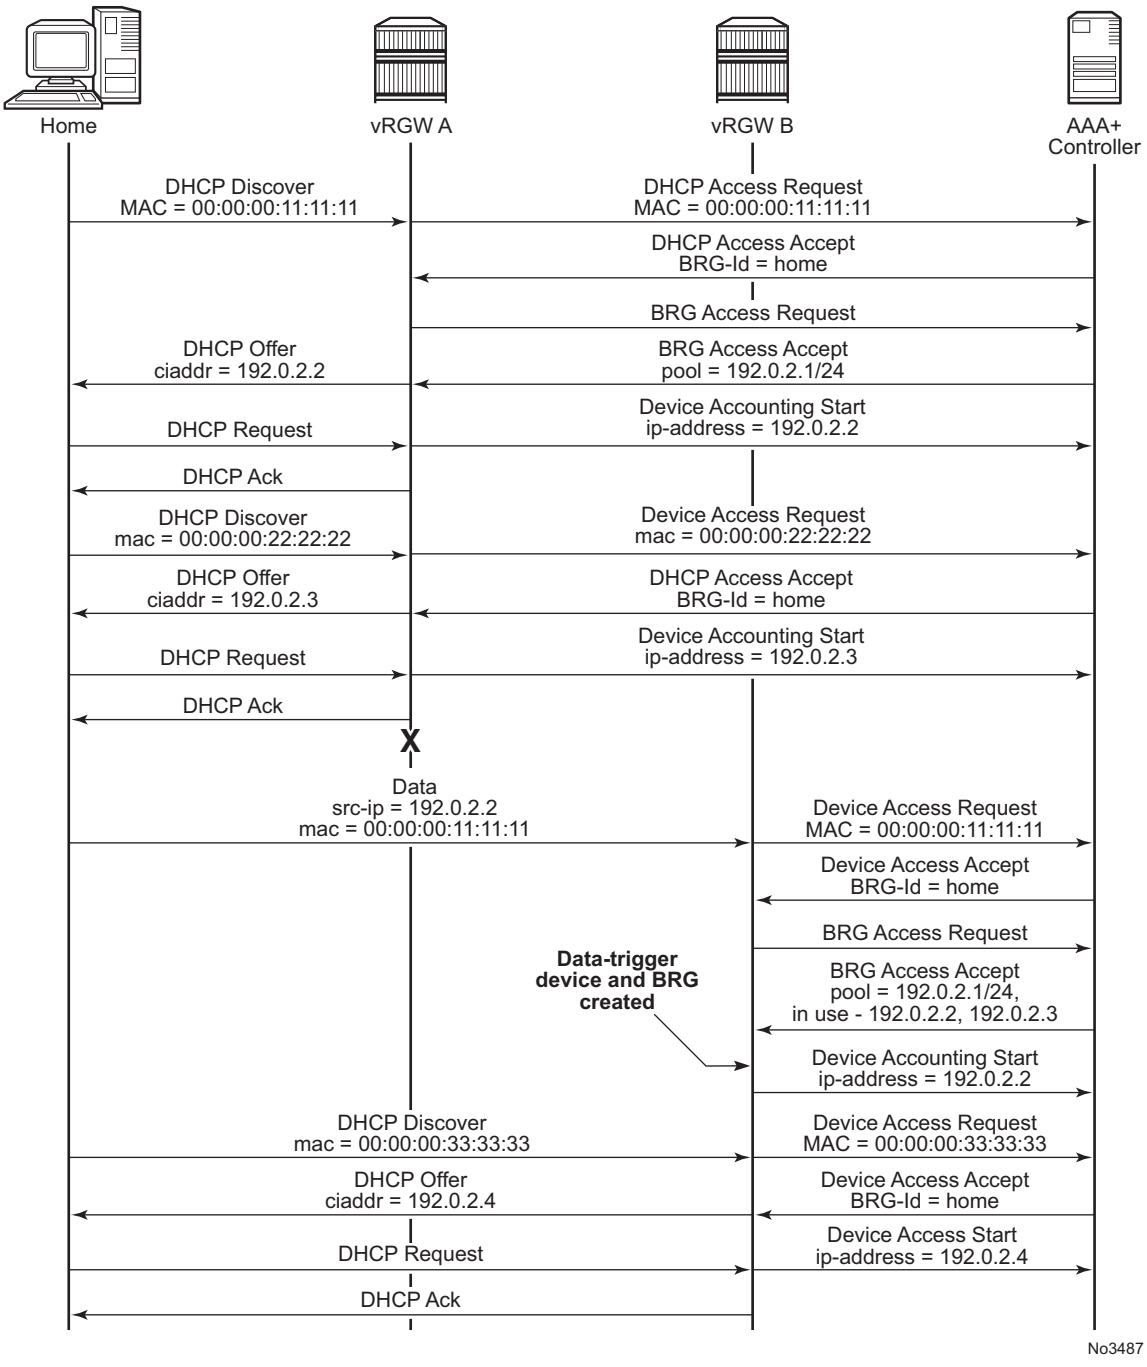

13.1.10.1 Pool State Synchronization

Because there is no stateful synchronization between pairs of redundant gateways, all pool state information is lost in a switchover event. The new gateway creates a new pool with state information for all configured reserved and static address, but has no information about dynamic addresses. When a data-triggered host is authenticated, the pool tries to recreate the lease corresponding with this source IP address. However, this is subject to race conditions. If a data trigger is delayed (for example, the device is not active upon switchover) the DHCP pool might already have given out the IP address to a new device in the home. This not only blocks the data-trigger for the old device, but might lead to address duplication issues inside the home.

To solve this issue, the controller keeps track of allocated dynamic IP addresses. After a switchover event, these addresses can be sent to the vRGW as part of the BRG configuration. In this case, the per-home pool only allocates these addresses to data-triggered devices or DHCP renew triggered devices that request this specific IP address. New devices receive addresses not present on this list. After a configurable timeout, the pool relinquishes any unclaimed IP addresses back to the pool for general use. The currently unclaimed addresses can be displayed using the following command.

```
show subscriber-mgmt brg gateway brg-id "00:00:5e:00:53:05" standby-ip-addresses
===============================================================================
Bridged Residential Gateway home-aware pool standby IP addresses
===============================================================================
Home-aware pool : 00:00:5e:00:53:05
-------------------------------------------------------------------------------
192.0.2.10
192.0.2.11
192.0.2.12
-------------------------------------------------------------------------------
No. of standby IP addresses: 3
===============================================================================
```

It is recommended to keep the pool state on the controller basic, by marking an address as used when a device is connected and unused when it is disconnected. Periodically, explicitly synchronize the controller with the vRGW to remove addresses for which a device disconnect notification may have been lost.

13.1.10.2 Regular Group Interfaces

On regular group interfaces, redundancy is supported by enabling SRRP and IPoE session stateless redundancy. For more information on stateless redundancy, see Stateful Multi-Chassis Redundancy (MCS). The backup router is always in the backup-routing mode. There is no support for shunting traffic via a redundant interface from the backup to master router.

IPv6 traffic and public IPv4 is attracted to the correct master router via the regular track-SRRP mechanism. L2-Aware NAT outside pools cannot use track-SRRP and should be configured distinctly on both routers making up the redundant pair. After a redundancy event, each home gains a new outside IP address on the new master router.

BRGs send out Gratuitous ARP (GARP) messages, which direct traffic to the correct master router. These GARPs update the FDB of the dual-homed Layer 2 aggregation nodes. GARPs are not supported for L2-Aware inside prefixes. At least one subscriber interface subnet must be explicitly configured to send a GARP.

When using MSAP deployments, the managed SAPs are not synchronized and must be recreated on the new master router via data-triggered authentication. Because there are no managed SAPs present, no GARPs can be sent when a backup router becomes master. If the aggregation node has a shared FDB for all CVLANs (FDB per SVLAN) or for all SVLANs (FDB per switch), it is recommended to configure a single static SAP per SVLAN or per port over which a single ARP can be sent.

13.1.10.3 WLAN-GW Group Interfaces

Redundancy on WLAN GW group interfaces reuses the WLAN-GW monitor route mechanism to determine the active or standby state of each router in a pair of redundant routers. For more information see WLAN-GW 1:1 Active-Backup Redundancy. When using the BRG MAC as a BRG identifier, an ARPoGRE (IPv4 GRE tunnel) or NDoGRE (IPv6 GRE tunnel) is triggered when receiving a data-trigger for an unknown device. Authentication of the data trigger is delayed until the ARP/ND is answered or timed out.

```
config>service>vprn>sub-if>grp-if>wlan-gw
        learn-ap-mac delay-auth
```

IPv6 and public IPv4 traffic from the network is attracted to the active router because a subscriber interface on a standby router is operationally disabled and its associated subnets are not exported to the route table. L2-Aware NAT outside pools stay active on both routers and should be configured independently on both routers, making up the redundant pair. After a redundancy event, each home gains a new outside IP address on the new master router.

Traffic from the BRG is directed by regular routing, as only the active router exports the configured tunnel endpoints.

Access via the Layer 2 access point is not supported in a redundant setup.

Alternatively, vRGW with tunnel access also supports active and active redundancy using two different tunnel endpoint addresses. In this setup, each vRGW acts independently and does not synchronize any state. Each BRG is responsible for determining which vRGW is considered active; for example, by sending ICMP/ICMP6 echo requests toward the tunnel endpoints. After detecting a failover, the BRG is direct traffic toward the alternate tunnel endpoint. The state on this vRGW is re-created using the data triggers. This is similar to the monitor-route active or backup redundancy as discussed in WLAN-GW Group Interfaces. Explicit BRG authentication is supported in this use case if the BRG keeps track of a pair of tunnel endpoints and RADIUS server (proxy) addresses.

Redundancy of public IPv4 addresses and IPv6 addresses is not supported with an active/active deployment. After a failover event, it is not possible to send traffic for these addresses until the BRG restores connectivity to its original vRGW.

13.1.11 BRG and vRG Caveats

- Static IPv6 hosts are not supported in vRGW. It is possible to provide a static IPv4 host with a SLAAC prefix and use IPoE linking to automatically create an associated IPv6 host.
- Wholesale/retail is not supported in vRGW.
- Subnets provisioned for BRG pool management must lie in a pre-configured L2-Aware NAT inside prefix. The dynamic range of a BRG pool must not contain the configured L2-Aware NAT inside IP address.
- BRG connectivity verification is limited to a maximum of 50,000 BRGs. If more BRGs are connected, connectivity verification is performed for the excess numbers and only hold-time applies before the BRG instance is deleted from the vRGW.
- On regular group interfaces, only a single BRG is supported per SAP.
- There is a maximum of one SLAAC prefix per BRG.

- There is a maximum of 128 IPoE sessions per BRG, and a maximum of 8 static hosts per BRG.
- The idle timeout is based on SLA profile instance, not per host. For hosts under the same BRG sharing an SLA profile, it is not possible to detect early disconnect of a single host.
- All SLAAC hosts under a BRG sharing the same prefix uses a common forwarding context downstream. For predictable behavior, the same SLA profile should be used for each SLAAC host.

13.1.12 External Allocation of L2-Aware NAT Outside IP Addresses

Certain residential deployments use multiple NAT outside IP addresses on routed physical raw (unknown encapsulation) packets, typically one per service. This feature provides similar support for multiple NAT outside IP addresses (up to four) per BRG (in other words, per subscriber) on vRGW. These addresses fall within locally-defined NAT pools on vRGW, but are assigned and managed by an external backend system. Each NAT outside IP address typically corresponds to a service, for instance, an HSI service may be NAT'ed to a different outside IP address than the voice service.

The NAT outside IP address and the corresponding NAT policy associated with the subscriber is provided by a VSA (Alc-Nat-Outside-IPs) in the RADIUS access-accept message or the CoA. Up to four instances of this VSA can be included in the RADIUS access-accept message or the CoA, which provides multiple NAT outside IP addresses for a BRG. The NAT pool referred to within the NAT policy must be configured for external assignment. If the provided outside IP address in the VSA does not fall in the NAT pool referenced within the corresponding NAT policy, then the outside IP address (and the mapping) is ignored. If the outside IP address falls within the NAT pool, and is not already allocated, then it is assigned to the L2-Aware subscriber. If the NAT policy contained in the VSA refers to a NAT pool that is not configured for external assignment, then the host setup fails. The external system is responsible for ensuring the stickiness of the outside IP addresses for the subscriber, if needed.

The system internally chooses an ISA in the NAT group to anchor an L2-Aware subscriber, such as the ISA, to where the NAT flows for the subscriber are created, based on upstream data or static port forwards. If a NAT outside IP address does not fall within the address block owned by the NAT ISA that anchors the L2-Aware subscriber, then the NAT outside IP address (/32) is added to the FDB with the ISA as the next hop. Otherwise, the downstream traffic forwarding to the NAT ISA for the L2-Aware subscriber follows the aggregate route corresponding to the address block owned by the NAT ISA.

The NAT outside IP address to use for a flow on the NAT ISA is based on a destination IP address lookup in a **nat-prefix-list** specified in the **sub-profile** for the subscriber. The **nat-prefix-list** contains a list of IP prefixes to NAT-policy mappings. The destination IP lookup in the **nat-prefix-list** provides the NAT policy to use. The NAT outside IP address that is associated with this NAT-policy is then used as the translated source IP.

```
config>service
     nat
       nat-policy "nat-policy-voice" create
           pool "voice-pool" router 401
       exit
       nat-prefix-list "nat-prefix-list-voice" application l2-aware-dest-to-
policy create
           prefix 203.0.113.235/16 nat-policy "nat-policy-voice"
       exit
     exit
config>subscr-mgmt
        sub-profile "sub-prof-hsi-voice" create
            nat-policy "default-policy"
            nat-prefix-list "nat-prefix-list-voice"
        exit
config>service>vprn>nat>outside
    pool "voice-pool" nat-group 1 type l2-aware create
    port-reservation blocks 1
    external-assignment
    exit
```

To add a new NAT outside IP address for a subscriber via CoA, AAA must provide updated <NAT outside IP-to-nat-policy> mappings. AAA associates the correct **sub-profile** (containing the **nat-prefix-list** to map a destination IP prefix with the NAT policy that contains the new NAT outside IP address) with the subscriber; that is, the sub-profile associated with the subscriber needs to be changed using CoA. The CoA must always contain all of the <NAT outside IP to-nat-policy> mappings associated with the subscriber (because it is cumulative). When the entire port range is available to the subscriber, the **port-reservation blocks** *num-blocks* command should be configured as 1 if **external-assignment** is enabled on the NAT pool.

All associated NAT outside IP addresses and corresponding NAT policies can be displayed via the **show service nat l2-aware-subscribers** command.

```
A:Dut-A# show service nat l2-aware-subscribers
```

```
===============================================================================
Layer-2-Aware NAT subscribers
===============================================================================
Subscriber                          : sub-2-4-ext
-------------------------------------------------------------------------------
ISA NAT group                       : 1
ISA NAT group member                : 1
UPnP policy                         : (None)
Default NAT policy                  : nat-policy-hsi
Per-host port block size            : N/A
Firewall policy                     : (None)

Policy                              : nat-policy-hsi
Purpose                             : nat
Outside router                      : vprn100
Outside IP                          : 198.51.100.235
DNAT default IP address override    : (Not Specified)
DNAT disabled by override           : false
Ports                               : 1024-5119

Policy                              : nat-policy-voice
Purpose                             : nat
Outside router                      : vprn100
Outside IP                          : 198.51.100.245
DNAT default IP address override    : (Not Specified)
DNAT disabled by override           : false
Ports                               : 1024-5119

-------------------------------------------------------------------------------
No. of subscribers: 1
```

The **nat pool** show command output shows the attribute that controls external assignment.

```
show router 70 nat pool "vprn l2aw"
===============================================================================
NAT Pool vprn l2aw
===============================================================================
Description                         : (Not Specified)
ISA NAT Group                       : 4
Pool type                           : l2Aware
Applications                        : (None)
Admin state                         : outOfService
Mode                                : auto (napt)
Port forwarding dyn blocks reserved : 0
Port forwarding range               : 1 - 1023
Port reservation                    : 128 blocks
Block usage High Watermark (%)      : (Not Specified)
Block usage Low Watermark (%)       : (Not Specified)
Block usage (%)                     : < 1
External assignment                 : true
Last Mgmt Change                    : 05/17/2016 13:41:04
===============================================================================
```

13.1.13 PPPoE Client

The vRGW allows operators to start a PPPoE client per BRG to a remote BNG. This PPPoE client forwards traffic that goes through NAT or firewall. The NAT-outside IP address and SLAAC prefix used are negotiated via the PPPoE client.

The PPPoE client runs directly on the ISA and encapsulates packets directly after NAT or firewall functionality. Encapsulated packets are sent out over an Epipe that is directly linked to the NAT group during configuration. Multiple clients can share a single Epipe if they have unique MAC addresses. Clients are automatically hashed over multiple ISAs in the NAT-group. By default, hashing is based on source and destination MAC; however, if many Epipes are being used and there is a high reuse of MAC addresses, Epipes should be configured to use per-service hashing instead.

A vRGW with PPPoE client enabled can still support routed hosts outside NAT or firewall.

13.1.13.1 PPPoE Client Setup

PPPoE client setup is triggered when a BRG is created with any PPPoE parameter provided during authentication. At least one PPPoE Epipe service ID and one PPPoE client policy name must be provided. A MAC address for the client should also be provided, but, if absent, the system attempts to parse the BRG ID as a MAC address. An optional service name can be configured that are reflected in PPPoE setup. LCP keepalives are defined in the PPPoE client policy. The client supports either no authentication, PAP, or CHAP. The username and password can be transmitted as parameters during BRG authentication.

The PPPoE client allows negotiation of both IPCP and IPv6CP with SLAAC and, optionally, stateless DHCPv6. The negotiated IPv4 address is used as a NAT-outside address and is not signaled to the clients in the home. The received IPv6 SLAAC prefix is reused as a SLAAC prefix towards the home. The Preferred and Valid lifetimes signaled towards home are not synchronized with the values received with the BRAS, but are only subject to local configuration. DHCPv6 hosts in the home cannot be subject to the PPPoE client.

The BRG can use the PPPoE client for address assignments instead of traditional sources such as AAA, LAA, or NAT pool. PPPoE client is mutually exclusive with other address sources. The used NAT and firewall policies must be configured in **l2-outside** mode. DNS, DNSv6, and NBNS can be provided by both the PPPoE client (IPCP, RA, stateless DHCPv6) or by vRGW AAA. Nokia recommends operators do not mix sources, because there is no strict precedence and only the last update is kept.

Setup of the BRG and devices linked to the BRG is blocked until setup of the PPPoE client is complete. If, after a timeout (configurable in the BRG profile) the client does not completely come up, setup of the BRG and devices continues with the information available. This allows, for example, IPv4 to continue setup even if IPv6 negotiation fails.

Aside from the MAC and service ID, most PPPoE client parameters can be changed while BRG is operational. Changing any PPPoE parameters causes the PPPoE client to restart, which impacts data traffic in the same way as a PPPoE client failure.

Non-LCP PPPoE control plane packets are always sent within the NC forwarding class and dot1p 0. Egress remarking can be used to change dot1p values.

13.1.13.2 PPPoE Client Failure

If there is a partial failure of the PPPoE client, no BRG or in-home device states are removed. All affected stacks no longer forwards traffic and any flow states may be removed. Setup of IPv4 hosts and routed (non-PPPoE client) devices continues. Existing SLAAC hosts are not removed, but a deprecation RA (with 0 lifetimes) is sent and no new SLAAC hosts are created.

The PPPoE client attempts to reestablish the connection, subject to the configured retry and backoff timers. Traffic is allowed again after the connection is restored. If a new IPv4 address is assigned, this is immediately applied as the new NAT outside address. If a new IPv6 prefix is assigned, IPv6 prefix replacement is triggered for existing SLAAC hosts.

If the peer MAC address changes during a reconnection, it may be hashed to a different NAT ISA. This is not supported by the system, and if it occurs, the entire BRG is removed.

13.1.13.3 LCP Keepalive

LCP keepalive intervals are configured in the PPPoE client policy. While the regular PPPoE control plane is initiated from the CPM, keepalives are handled directly on the ISA. Because the ISA is aware of any traffic, this handling is optimized so that echo requests are not sent as long as active traffic is received from the BRAS. The keepalive timer is reset upon each received echo-request. These two mechanisms significantly reduce PPP keepalive messages without compromising liveness detection. Incoming echo requests are always responded to, as long as LCP is in the LCP Opened state.

LCP PPPoE packets are always sent within the NC forwarding class and dot1p 7. Egress remarking can be used to change dot1p values.

13.1.13.4 MRU/MTU

Each PPPoE client policy allows configuration of both MRU and MTU. The MRU is sent towards the peer during the discovery phase. If any of the MRU or MTU is bigger than 1492, the Max-Payload tag is included and is set to the maximum of the configured MRU and MTU values.

Upstream data packets are subject to the smallest of either the received MRU or the configured MTU. The node fragments IPv4 packets without the DF flag set and drops IPv6 packets and IPv4 packets with the DF flag set. For the packets with the DF flag set, appropriate ICMP or ICMP v6 error messages are sent.

13.1.14 SLAAC Prefix Replacement

SLAAC prefix deprecation is supported with vRGW. Contrary to generic prefix replacement, vRGW uses radius attribute Framed-IPv6-Prefix to change the prefix in a BRG-level CoA or re-authentication. For PPPoE clients, prefix replacement is performed whenever the SLAAC prefix is signaled via PPPoE changes. See SLAAC for more information about SLAAC prefix deprecation.

When a SLAAC prefix is assigned to a BRG that previously had no prefix, an unsolicited RA is generated for all sessions currently active for that BRG. Data-triggered host creation should be enabled to allow subsequent setup of IPv6 hosts.

SLAAC prefix deprecation is supported for a PPPoE client. Whenever the state of the SLAAC prefix is lost, such as when a PPPoE client goes down, an RA with 0 lifetimes is sent out and no new SLAAC hosts can be set up. Existing hosts remain in the system but are no longer forwarded. When the same SLAAC prefix is reassigned, existing hosts begin forwarding traffic again.

13.2 Home LAN Extension

This section describes the Home LAN Extension (HLE) feature.

13.2.1 Overview

HLE allows an operator to extend the home network of the broadband user to a WAN network, such as a data center, by creating a per-home BD (BD) on a Broadband Network Gateway (BNG). This feature provides the operator with the capability to deploy new services in a data center that have full Layer 2 reachability to the home and visibility to each individual host.

The per-home BD is created on a WLAN-GW ISA, and the system uses standard BGP EVPN VPLS services to extend the BD to remote networks.

Figure 198 displays an example of an HLE configuration.

Figure 198 **Home LAN Extension**

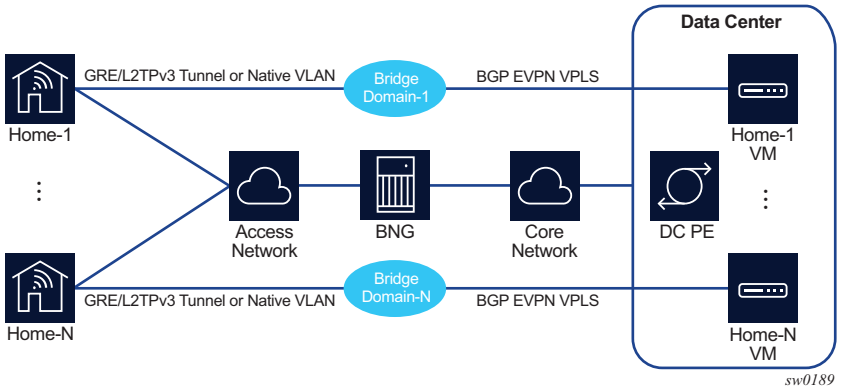

HLE requires the BNG to have Layer 2 access to the home network. The BNG supports the following types of access:

- soft GRE or L2TPv3 tunnel from a home gateway, which encapsulates Ethernet traffic from the host into a GRE/L2TPv3 tunnel
- native VLAN access that is terminated on the WLAN-GW group ISA using L2-AP access

A unique BD is created on the ISA for each home. The BD bridges traffic between the following connections:

- access-facing connection (for example, home) — GRE/L2TPv3/L2-AP
- network-facing connection (for example, DC) — BGP EVPN tunnel
- ESM SAP-facing connection — each home has its own ESM SAP

Each BD has a unique ID, which is a number returned by the RADIUS server as the Alc-Bridge-Id attribute during home authentication.

With HLE services, each home host is a WLAN-GW UE object and also an ESM host object. Each network host (such as a VM in a data center) is a WLAN-GW UE object but not a ESM host. This means that a network host cannot use the BNG as the default router for other non-home-facing traffic.

HLE relies on SR OS vRGW functions, which means that the vRGW BRG in the same VLAN range of the WLAN-GW group interface must be enabled for HLE.

13.2.2 Authentication and Authorization

Because HLE relies on the vRGW, it is carried in conjunction with vRGW home authentication and supports both implicit and explicit vRGW home authentication. Enabling the HLE service for a specific home (ESM subscriber) is triggered by a RADIUS server that returns an Alc-Bridge-Id message during home authentication.

The Alc-Bridge-Id attribute must be returned in both home and session authentication. The returned value must be equal on both authentication levels.

The Alc-Bridge-Id is the ID of the per-subscriber BD on the ISA. It is different from the Alc-BRG-ID, which is the ID of the BRG.

In addition to the Alc-Bridge-Id, the following optional attributes can be returned during BRG authentication:

- HLE BGP EVPN route target: Alc-RT
- HLE BGP EVPN route distinguisher: Alc-RD
- HLE BGP EVPN VXLAN VNI: Alc-Vxlan-VNI

13.2.3 Data Plane Tables

HLE is an essential Layer 2 bridge service. The data plane consists of a per-subscriber BD that contains the following tables:

- MAC table — this table includes the learned MAC address by access and network connections
- flood table — this table includes flood destinations for broadcast, unknown unicast, multicast (BUM) traffic and typically contains access, and network connections
- IPv4 ARP table — this table includes learned ARP entries by access and network connections
- IPv6 neighbor table — this table includes learned neighbor entries by access and network connections

The ARP and neighbor tables are populated when Assistive Address Resolution (AAR) is enabled.

When MAT is disabled, the ISA forwards packets based on the destination MAC address by performing a lookup in the MAC table and floods BUM traffic to destinations in the flood table.

For flooding traffic, the system does not send a to-be-flooded packet back to where it was received. If the packet is received from one remote network connection, the system does not forward it to another remote network connection.

13.2.4 BGP EVPN VPLS

HLE uses standard BGP EVPN VPLS services to extend the BD to the remote network. A BGP EVPN VPLS service is automatically created when a BD is created on the ISA.

HLE only supports VXLAN as a BGP EVPN data plane. When the next-hop field of the MP_REACH_NLRI is the VXLAN VTEP, it supports the following BGP EVPN route types:

- Type 2 MAC or IP advertisement route:
  - ESI: 0
  - Ethernet-tag: 0
  - IP address field:
    - sending: not used
    - receiving: used to populate the BD’s ARP/neighbor table
  - Label-1: VNI
    Label-2: not used
  - BGP encapsulation extended community with tunnel type set to VXLAN

  - MAC mobility extended community: seq=0, sticky=1 (only when MAC NAT is enabled)
- Type 3 inclusive multicast Ethernet tag route:
  - originating router address: system interface address
- BGP encapsulation extended community route with tunnel type set to VXLAN

PMSI tunnel attributes include:

- flags: 0
- tunnel type: ingress replication
- MPLS label:
  - sending: local VNI
  - receiving: remote VNI
- tunnel ID: VXLAN VTEP
  - sending: VTEP of the corresponding ISA
  - receiving: remote flooding destination

The VXLAN packet is forwarded in the base routing instance

13.2.5 Assistive Address Resolution

Assistive Address Resolution (AAR) is an optional HLE feature used to avoid sending ARP/ND requests from the home across a WAN to a remote network, or sending ARP/ND requests from the remote network across an access network to a home. The system responds to the ARP/ND request from the network or home with the learned ARP/ND entries, instead of flooding it.

With AAR, the system populates the ARP and neighbor tables with the learned ARP and neighbor entries, using the following methods:

- BGP EVPN MAC routes that contain an IP address
- (G)ARP/ND/NS packets

When the ISA receives an ARP/ND request, it performs a lookup in the BD’s ARP and neighbor tables by using a target IP address as the key. The following conditions apply:

- If the requested IP address is over a different connection, then the ISA generates an ARP/ND reply with the match result without flooding the request.

- If the requested IP address is over same connection, then the ISA drops the request. If there is no match, then the ISA floods the address.

13.2.6 MAC Address Translation

MAC address translation (MAT) is an optional HLE feature that translates the access or network host's MAC address into a single BD MAC address. This feature decreases the number of BGP EVPN MAC routes-to-advertise per subscriber to one and also eliminates the BGP update when hosts are created and removed which increases BGP's stability.

MAT is performed on traffic in both directions:

- traffic from access to network — the system changes the source MAC address to BD MAC address before forwarding the packet to the network direction and change the destination MAC address into a real MAC address of the network host based on a ARP table lookup by using the destination IP address as key.
- traffic from network to access — the system changes the destination MAC address (BD MAC) to the real host MAC address based on the ARP table or neighbor table lookup using the IP address as key, and change the source MAC address to a BD MAC address before forwarding the packet to the access direction.

With MAT enabled, upon receiving an ARP or neighbor request, the system performs a lookup in the ARP or neighbor table with the target IP address which determines the next action:

- when the target IP address is known:
  - the system responds with a BD MAC address
- when the target IP address is unknown, the request is flooded and the ARP or neighbor table is populated with the response.

Because of the ARP processing, the destination MAC of the received unicast packets would be BD MAC. The system only advertises BD MAC in EVPN when MAT is enabled. For unicast data packets received, if the destination MAC address is a BD MAC address, the ISA performs a lookup in the ARP or neighbor table by using the target IP address as key.

- If there is a match, the destination MAC address is changed to the MAC address that resulted from the ARP or neighbor table lookup and change the source MAC address to the BD MAC.
- If there is no match if it is unicast ARP/ND packet, then it is flooded, otherwise, it is forwarded to the BD's ESM SAP.

If the destination MAC address is a unicast MAC address, but not a BD MAC address, the destination MAC address is forwarded based on the MAC table lookup.

If the destination MAC address is a broadcast or multicast MAC address, then the packet is flooded.

MAT requires assistive address resolution to be enabled and with MAT enabled, the system also changes the source and target hardware addresses in ARP/ND requests and replies.

13.2.7 Configuring HLE

The following is an example of a configuration to enable HLE in the system with implicit home authentication.

1. Configure the vRGW BRG.
2. Enable HLE in the base routing instance:

```
config>router>vrgw# info
----------------------------------------------
            lanext
                vxlan-vtep-range start 198.51.100.235 end 198.51.100.245
                wlan-gw-group 1
                no shutdown
            exit
----------------------------------------------
```

3. Configure the HLE EVPN route target number:

```
config>subscr-mgmt>vrgw>lanext# info
----------------------------------------------
                router-target-as-number 100
----------------------------------------------
```

4. Configure the maximum number of BDs under the group interface:

```
config>service>vprn>sub-if>grp-if>wlan-gw
                max-lanext-bd 100
```

5. Configure BGP in the base routing instance:

```
config>router>bgp# info
----------------------------------------------
            group "evpn"
                family evpn
                type internal
                neighbor 2.2.2.2
                exit
            exit
            no shutdown
----------------------------------------------
```

6. Configure HLE under the VLAN range:

```
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range# info
----------------------------------------------
          authentication
              authentication-policy "isa-rad"
          exit
          authenticate-on-dhcp
          vrgw
              brg
                  default-brg-profile "brgp-1"
                  no shutdown
              exit
              lanext
                  assistive-address-resolution
                  bd-mac-prefix AA:BB:CC
                  mac-translation
                  no shutdown
              exit
          exit
----------------------------------------------
```

7. Provision the RADIUS server to include Alc-Bridge-Id in the host and BRG records.

13.3 AP Agnostic Access for Multiple Dwelling Units

13.3.1 Overview

In a typical vRGW deployment, including HLE, the subscriber's BRG instance and BD on vRGW are tied to an access circuit (such as a soft GRE or soft L2TPv3 tunnel) from a single bridged Access Point (AP) or a residential gateway (RG). This feature adds support for subscriber to be AP agnostic. This means that a subscriber's BRG instance and BD are not tied to a single bridged AP or RG.This is particularly useful when the customer premise is an multi-dwelling (MDU) unit inhabited by multiple independent tenants where these tenants within the building can obtain connectivity from any bridged AP in the building. Bridged WIFI AP and RGs can be installed in various parts of the building and are not owned or operated by specific tenants. Each AP can be provisioned with a common SSID (for example, an operator-branded SSID providing bulk Internet access and intra-MDU connectivity). Each AP accesses the network by an L2oGRE or L2TPv3 tunnel terminating on a gateway (vRGW) that provides integrated bridging and vRGW processing. The existing vRGW functionality is defined in Virtual Residential Gateway. The per tenant (access) bridging function on vRGW is described in Home LAN Extension. With this AP agnostic access feature, the traffic flow is handled as follows (Figure 199).

- Traffic between two devices of a tenant behind the same AP is locally switched by the AP.
- Traffic between two devices belonging to the same tenant connecting from two different APs is not locally switched by the APs but is tunneled to the vRGW or gateway that provides a bridged domain per tenant. This traffic is bridged by the gateway using the per-tenant bridge-domain (BD).
- Traffic between two devices belonging to different tenants connecting from the same or different AP in the building are tunneled to the gateway or vRGW. This traffic is subject to vRGW processing, such as ESM followed by L2-aware NAT. The isolation between tenants is provided by separate BRG contexts per tenant on the vRGW.
- Traffic from a tenant device, connecting from any AP in the building, that is destined to any destination on the Internet is forwarded by the AP over the tunnel and is subject to vRGW processing; for example, ESM followed by NAT.

Figure 199 **AP Agnostic Access – Integrated Bridging and vRGW Processing**

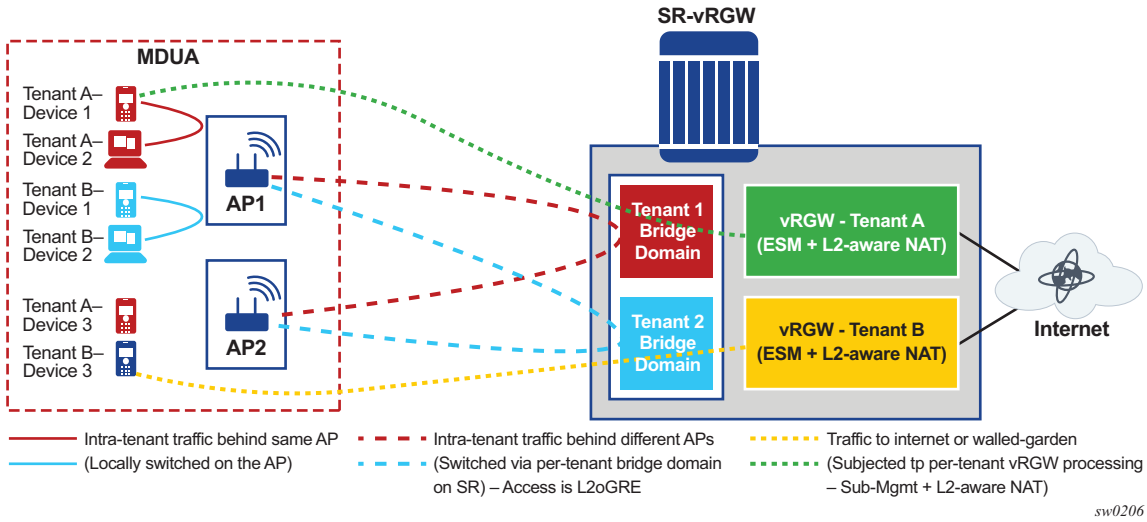

13.3.2 Bridge Domain and BRG Identification

The tenant's bridge domain (BD) and BRG instance (for vRGW processing) is not identified by the tunnel source's IP address. Each tenant is provided a unique dot1q tag, which is maintained by the back-end system. The traffic from a device that is tunneled to vRGW carries this unique dot1q tag. The AP is aware of the tag per tenant and can provide traffic isolation between tenants.The BD and BRG instance for the tenant on vRGW are identified by the combination of the dot1q tag and the tunnel destination (that is, a gateway address configured under the WLAN-GW group-interface). Therefore, devices belonging to a tenant can connect from any AP, as the tenant context is not identified based on tunnel source (that is, the AP's IP address). This AP agnostic connectivity can be enabled on a per-VLAN range basis (in the **vrgw>lanext** context) as shown in the following example.

```
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>
     vrgw
       lanext
          access
            [no] multi-access
          exit
       exit
     exit
```

The **multi-access** keyword indicates that the connecting device is not tied to a single AP and implicitly enables the identification of the BD and BRG instance for this device based on dot1q tag in the frame, and the gateway address, such as the L2oGRE or L2TPv3 tunnel endpoint IP address.

13.3.3 ARP Handling

Split-horizon semantics are maintained for ARP handling with the per-tenant BD and ARP table. If the target IP of an incoming ARP request on the tunnel is in the ARP table for the tenant, and the tunnel source of the ARP request matches the tunnel source associated with the target IP, then no ARP response is generated. A lack of ARP response implies that both the source and target are behind the same AP and the ARP is handled on a local WLAN on the AP. If the target IP in the ARP request is not found in the tenant's ARP table, then the ARP request is flooded to all the access tunnels, such as to all known tunnel sources for the MAC addresses in the bridge-domain.

13.3.4 Mobility

An existing UE belonging to a tenant can connect from a different AP than its initial connection. This can be either due to mobility or a change in the signal strength of those APs in range. Both control and data-triggered mobility within context of the dot1q tag are supported. Control mobility includes DHCPv4, DHCPv6, and RS messages. All UEs for a given tenant are anchored on the same ISA. When mobility is detected (by a change in the tunnel source for the UE), the impacted UE MAC address in the BD is updated to point to the new tunnel. Triggered interim accounting updates are generated, if enabled.

vRGW functions in the SR OS are supported for this AP agnostic access.

13.4 vRGW and HLE Command Reference

13.4.1 Command Hierarchies

- Subscriber Management BRG Profile Commands
- Subscriber Management Commands
- Service Commands
- Router HLE Commands
- Subscriber Management HLE Commands
- Service Commands
- Show Commands
- Clear Commands
- Tools Commands

13.4.1.1 Subscriber Management BRG Profile Commands

```
configure
    — subscriber-mgmt
        — vrgw
            — brg
                — brg-profile profile-name [create]
                — no brg-profile profile-name
                    — connectivity-verification [count nr-of-attempts] [timeout
                          timeout-seconds] [retry-time retry-seconds]
                    — no connectivity-verification
                    — description description-string
                    — no description
                    — dhcp-pool
                        — lease-time seconds
                        — no lease-time
                        — options
                            — custom-option option-number address [ip-
                                  address]
                            — custom-option option-number hex hex-string
                            — custom-option option-number string ascii-string
                            — no custom-option option-number
                            — subnet ip-prefix/prefix-length start ip-address
                                  end ip-address
                        — standby-ip-lifetime [days days] [hrs hrs] [min min] [sec
                              sec]
```

```
            — standby-ip-lifetime [seconds]
            — no standby-ip-lifetime
        — hold-time seconds
        — no hold-time
        — initial-hold-time seconds
        — no initial-hold-time
        — radius-authentication
            — password password [hash | hash2]
            — no password
            — radius-server-policy policy-name
            — no radius-server-policy
        — [no] radius-proxy-server router router-instance name server-
                name
        — sla-profile-string string
        — no sla-profile-string
        — sub-profile-string string
        — no sub-profile-string
        — uplink-initial-wait seconds
```

13.4.1.2 Subscriber Management Commands

```
configure
    — subscriber-mgmt
        — authentication-policy name [create]
        — no authentication-policy name
            — [no] include-radius-attribute
                — [no] wifi-num-attached-ues
        — pppoe-client-policy pppoe-client-policy-name [create]
        — no pppoe-client-policy pppoe-client-policy-name
            — description description-string
            — no description
            — keepalive seconds [hold-up-multiplier multiplier]
            — no keepalive
            — mru mru-bytes
            — no mru
            — mtu mtu-bytes
            — no mtu
            — python-policy name
            — no python-policy
            — restart-backoff initial-time seconds max-time seconds
            — no restart-backoff
            — stack [ipv4] [ipv6-slaac]
            — no stack
        — radius-accounting-policy name [create]
        — no radius-accounting-policy name
            — [no] include-radius-attribute
                — [no] wifi-num-attached-ues
                — [no] brg-num-active-sessions
```

13.4.1.3 Service Commands

```
config>service
    — ies service-id|vprn service-id
        — subscriber-interface ip-int-name [create]
        — no subscriber-interface ip-int-name
            — group-interface ip-int-name [create]
            — no group-interface ip-int-name
                — brg
                    — [no] authenticated-brg-only
                    — default-brg-profile profile-name
                    — no default-brg-profile
                    — [no] shutdown
                — wlan-gw
                    — vlan-tag-ranges
                        — range start [value] end [value]
                        — range default
                        — no range start [value] end [value]
                            — [no] data-triggered-ue-creation
                            — vrgw
                                — brg
                                    — [no] authenticated-brg-only
                                    — default-brg-profile profile-name
                                    — no default-brg-profile
                                    — [no] shutdown
                            — xconnect
                                — accounting-policy isa-radius-policy-
                                    name
                                — no accounting-policy
                                — accounting-update-interval [interval]
                                — no accounting-update-interval
                                — [no] mobility-acct-updates
                                — [no] shutdown
                    — learn-ap-mac [delay-auth]
                    — no learn-ap-mac
```

13.4.1.4 Router HLE Commands

```
configure
    — router
        — vrgw
            — [no] lanext
                — [no] shutdown
                — vxlan-port [4789 | 8472]
                — no vxlan-port
                — vxlan-vtep-range start ip-address | ipv6-address end ip-address |
                    ipv6-address
                — no vxlan-vtep-range
                — wlan-gw-group wlan-gw-group-id
                — no wlan-gw-group
```

```
— wlan-gw
    — xconnect
        — [no] shutdown
        — tunnel-source-ip ipv6-address/prefix-length
        — no tunnel-source-ip
        — wlan-gw-group nat-group-id
        — no wlan-gw-group
```

13.4.1.5 Subscriber Management HLE Commands

```
configure
    — subscriber-mgmt
        — radius-accounting-policy name [create]
        — radius-accounting-policy
            — [no] include-radius-attribute
                — [no] lanext-bridge-id
                — [no] lanext-device-type
                — [no] lanext-route-distinguisher
                — [no] lanext-route-target
                — [no] lanext-vni
        — vrgw
            — lanext
                — router-target-as-number as-number
                — no router-target-as-number
```

13.4.1.6 Service Commands

```
config>service>ies service-id|vprn service-id
    — subscriber-interface ip-int-name [create]
    — no subscriber-interface ip-int-name
        — group-interface ip-int-name [create]
        — no group-interface ip-int-name
            — wlan-gw
                — max-lanext-bd [value]
                — no max-lanext-bd
                — vlan-tag-ranges
                    — range start [value] end [value]
                    — range default
                    — no range start [value] end [value]
                        — vrgw
                            — [no] lanext
                                — [no] access
                                    — max-mac [value]
                                    — no max-mac
                                    — [no] multi-access
                                — [no] assistive-address-resolution
                                — bd-mac-prefix mac-prefix
                                — no bd-mac-prefix
```

```
— [no] mac-translation
— [no] network
      — max-mac [value]
      — no max-mac
      — [no] shutdown
— [no] shutdown
```

13.4.1.7 Show Commands

```
show
   — subscriber-mgmt
      — vrgw
         — brg
            — brg-hosts [mac ieee-address] [ip ip-address]
            — brg-profile
            — brg-profile name
            — brg-profile name associations
            — gateway brg-id brg-ident
               — bindings [mac ieee-address]
               — host mac ieee-address ip ip-address
               — hosts
               — standby-ip-addresses
            — gateways
               — gateways
               — gateways brg-profile profile-name
            — pppoe-client [brg-id [32 chars max]]
            — pppoe-client [brg-id [32 chars max]] [detail]
            — pppoe-client [brg-id [32 chars max]] [statistics]
         — lanext
            — bd [bridge-id]
```

13.4.1.8 Debug Commands

```
debug
   — subscriber-mgmt
      — vrgw
         — brg
            — brg-id brg-ident
            — detail-level {low | medium | high}
            — direction {ingress | egress | both}
            — discovery [padi] [pado] [padr] [pads] [padt]
            — [no] dropped-only
            — ppp [lcp] [pap] [chap] [ipcp] [ipv6cp] [ipv6]
```

13.4.1.9 Clear Commands

```
clear
    — subscriber-mgmt
        — vrgw
            — brg
                — gateway brg-id brg-ident
                — gateway brg-id brg-ident idle-bindings [binding ieee-address]
                — gateway all-gateways
                — gateway brg-id brg-ident all-hosts
                — gateway brg-id brg-ident host ieee-address
```

13.4.1.10 Tools Commands

```
tools
    — dump
        — wlan-gw
            — lanext
                — bd bridge-id
                    — arp-table [ip ipv4-address] [class class] [next-index index]
                    — flood-table [class class] [detail]
                    — mac-table [mac ieee-address] [class class] [next-index index]
                      [detail]
                    — neighbor-table [ip6 ipv6-address] [class class] [next-index
                      index]
    — perform
        — subscriber-mgmt
            — vrgw
                — brg
                    — pppoe-client
                        — restart brg-ident
        — wlan-gw
            — lanext
                — bd bridge-id
                    — clear-arp [ip ipv4-address]
                    — clear-mac [mac ieee-address]
                    — clear-neighbor [ip6 ipv6-address]
```

13.4.2 Command Descriptions

13.4.2.1 Generic Commands

description

Syntax **description** *description-string*
no description

Context config>subscr-mgmt>pppoe-client-policy
config>subscr-mgmt>vrgw>brg>brg-profile
config>router>vrgw>lanext
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext>network
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext>network

Description This command creates a text description stored in the configuration file for a configuration context. This associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes the description string from the context.

Parameters *description-string* — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

shutdown

Syntax [**no**] **shutdown**

Context config>service>ies>sub-if>grp-if>wlan-gw>vlan-ranges>range>vrgw>brg
config>service>vprn>sub-if>grp-if>wlan-gw>vlan-ranges>range>vrgw>brg
config>service>ies>sub-if>grp-if>brg
config>service>vprn>sub-if>grp-if>brg

Description This command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they can be deleted.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system-generated configuration files.

The **no** form of this command puts an entity into the administratively enabled state.

13.4.2.2 vRGW Commands

vrgw

Syntax **vrgw**

Context config>subscr-mgmt
config>service>ies>sub-if>grp-if>wlan-gw>vlan-ranges>range

Description This command enables the context to configure Virtual Residential Gateway parameters.

brg

Syntax **brg**

Context config>subscr-mgmt>vrgw

Description This command enables the context to configure Bridged Residential Gateway parameters.

brg-profile

Syntax **brg-profile** *profile-name* [**create**]
no brg-profile *profile-name*

Context config>subscr-mgmt>vrgw>brg

Description This command creates the profile for Bridged Residential Gateway (BRG) devices. The BRG profile specifies default parameters that are used for host management under a single BRG.

The **no** form of this command removes the profile name from the configuration.

Parameters *profile-name* — Specifies the name of the BRG profile.

create — Keyword used to create a BRG profile. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

connectivity-verification

Syntax **connectivity-verification** [**count** *nr-of-attempts*] [**timeout** *timeout-seconds*] [**retry-time** *retry-seconds*]
no connectivity-verification

Context config>subscr-mgmt>vrgw>brg>brg-profile

Description This command configures the BRG connectivity verification. The system uses ICMP Echo request messages for connectivity verification.

When the last host associated with a BRG is removed, a ping mechanism is used to verify if the BRG is still active. This command specifies the parameters used in this mechanism.

The **no** form of this command disables the BRG ping mechanism and removes the BRG without verification. Any configured hold time still applies.

Default connectivity-verification count 3 timeout 30 retry-time 900

Parameters *nr-of-attempts* — Specifies the number of connectivity verification attempts this system makes before a BRG is considered down.

Values 1 to 5

timeout-seconds — Specifies the time, in seconds, after which an unanswered ping is considered failed.

Values 5 to 60

retry-seconds — Specifies the time, in seconds, that the system waits while it considers a BRG down before it starts a new connectivity verification cycle. If a ping succeeds, the mechanism will be retried after this time.

Values 300 to 3600

dhcp-pool

Syntax **dhcp-pool**

Context config>subscr-mgmt>vrgw>brg>brg-profile

Description This command enables the context to configure per-subscriber IPv4 address pool parameters to be used for address allocation. Pools for different subscribers can overlap. Specific pool parameters can be overridden by RADIUS.

lease-time

Syntax **lease-time** *seconds*
lease-time [**days** <*days*>] [**hrs** <*hrs*>] [**min** <*min*>] [**sec** <*sec*>]
no lease-time

Context config>subscr-mgmt>vrgw>brg>brg-profile>dhcp-pool

Description This command configures the lease time, in seconds, to be used when allocating addresses from the pool. This time value should always be longer than the renew/rebind time.

The **no** form of this command reverts to the default.

Default lease-time hrs 6

Parameters *seconds* — Specifies the lease time in seconds.

Values 300 to 315446399

days — Specifies the lease time in days.

Values 1 to 3650

hrs — Specifies the lease time in hours.

Values 1 to 23

min — Specifies the lease time in minutes.

Values 1 to 59

sec — Specifies the lease time in seconds.

Values 1 to 59

options

Syntax **options**

Context config>subscr-mgmt>vrgw>brg>brg-profile>dhcp-pool

Description This command enables the context to configure options that are reflected in DHCP.

custom-option

Syntax **custom-option** *option-number* **address** [*ip-address*]
custom-option *option-number* **hex** *hex-string*
custom-option *option-number* **string** *ascii-string*
no custom-option *option-number*

Context config>subscr-mgmt>vrgw>brg>brg-profile>dhcp-pool>options

Description This command configures DHCP options.

Parameters *option-number* — Specifies the number of this DHCP option.

ip-address — Specifies up to four IP addresses.

hex-string — Specifies a hex string.

ascii-string — Specifies the ASCII string.

subnet

Syntax **subnet** *ip-prefix/prefix-length* **start** *ip-address* **end** *ip-address*

Context config>subscr-mgmt>vrgw>brg>brg-profile>dhcp-pool

Description This command configures the subnet that will be used for the L2aware-subscriber. This subnet is only locally significant and can overlap with other subscribers. The subnet is derived by ignoring the host bits of the IP address. The IP address specifies the default gateway that will be signaled in DHCP along with the netmask derived from the prefix length.

The start and end addresses specify the range of addresses that are suitable for allocation within the given subnet. If the subnet address (host bits 0), broadcast address (host bits 1) or default-gw address fall in this range, they will not be considered for allocation.

Changing the subnet will only affect new subscribers. New and existing hosts for existing subscribers will keep allocating from the original subnet.

Default subnet 192.168.0.1/24 start 192.168.0.2 end 192.168.0.254

Parameters *ip-prefix/prefix-length* — Specifies the IP prefix and prefix length.

start *ip-address* — Specifies the starting IP address.

end *ip-address* — Specifies the ending IP address.

standby-ip-lifetime

Syntax **standby-ip-lifetime** [**days** *days*] [**hrs** *hrs*] [**min** *min*] [**sec** *sec*]
standby-ip-lifetime [*seconds*]
standby-ip-lifetime

Context config>subscr-mgmt>vrgw>brg>brg-profile>dhcp-pool

Description This command defines how long these addresses will be kept aside when standby addresses are signaled to the pool. During this time these addresses can only be used by devices explicitly requesting the IP (for example, datatrigger or DHCP renew/rebind). When the timer expires the addresses will again become available for dynamic allocation.

Default standby-ip-lifetime hrs 6

Parameters *days* — Specifies the standby IP lifetime in days.

Values 1 to 3650

hrs — Specifies the standby IP lifetime in hours.

Values 1 to 23

min — Specifies the standby IP lifetime in minutes.

Values 1 to 59

sec — Specifies the standby IP lifetime in seconds.

Values 1 to 59

seconds — Specifies the lifetime of the standby IP addresses

Values 300 to 315446399

seconds — Specifies the lifetime in seconds.

Values 300 to 315446399

hold-time

Syntax **hold-time** *seconds*
hold-time [**days** *days*] [**hrs** *hrs*] [**min** *min*] [**sec** *sec*]
no hold-time

Context config>subscr-mgmt>vrgw>brg>brg-profile

Description This command holds the BRG object for the specified time. This applies when the connectivity verification fails or when the last host is removed and **no connectivity-verification** is enabled. Hold time does not apply to an explicit removal via the **radius** or **clear** commands.

The **no** form of this command disables the hold time.

Parameters *seconds* — Specifies the time, in seconds, to hold on to a BRG after the system considered it down.

Values 30 to 2592000

days — Specifies the hold-time in days.

Values 1 to 30

hrs — Specifies the hold-time in hours.

Values 1 to 23

min — Specifies the hold-time in minutes.

Values 1 to 59

sec — Specifies the hold-time in seconds.

Values 1 to 59

initial-hold-time

Syntax **initial-hold-time** *seconds*

no initial-hold-time

Context config>subscr-mgmt>vrgw>brg>brg-profile

Description This hold time applies to BRG instances that are created without any hosts; for example, triggered due to proxy authentication. During this initial hold-time the BRG is not deleted and no connectivity-verification is started. When this timer expires, connectivity verification is started. If connectivity verification is disabled and no hosts are associated with the BRG upon expiry, the regular hold time is started and the BRG instance will be removed.

This command allows an operator to have fast BRG removal (no connectivity-verification and **no hold-time**) but still use BRG proxy authentication. Without an initial hold time the BRG would be removed immediately after creation.

This command does not apply to BRG instances that are created through host setup.

The no form of this command reverts to the default.

Default initial-hold-time min 5

Parameters *seconds* — Specifies the initial time, in seconds, to hold on to a BRG after the system considered it down.

Values 0 to 900

min — Specifies the initial time in minutes.

Values 1 to 15

sec — Specifies the initial time in seconds.

Values 1 to 59

radius-authentication

Syntax **radius-authentication**

Context config>subscr-mgmt>vrgw>brg>brg-profile

Description This command enables the context to configure parameters related to RADIUS authentication performed for the BRG.

password

Syntax **password** *password* [**hash** | **hash2**]
no password

Context config>subscr-mgmt>vrgw>brg>brg-profile>radius-authentication

Description This command configures the value of the user password as signaled in RADIUS.

The **no** form of this command indicates that no password field should be sent in RADIUS.

Parameters **hash** — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

radius-server-policy

Syntax **radius-server-policy** *policy-name*
no radius-server-policy

Context config>subscr-mgmt>vrgw>brg>brg-profile>radius-authentication

Description This command is used if the BRG needs to be authenticated to the controller/PCMP by the vRGW. This is required if the BRG does not perform RADIUS authentication via the proxy server. The vRGW will originate a valid Access Request using the BRG ID as the username.

The **no** form of this command removes the **radius-server-policy** from the configuration. Setup of an unauthenticated BRG will now fail.

Parameters *policy-name* — Specifies the RADIUS server policy, up to 32 characters, to be applied to this subscriber authentication policy.

radius-proxy-server

Syntax [**no**] **radius-proxy-server router** *router-instance* **name** *server-name*

Context config>subscr-mgmt>vrgw>brg>brg-profile

Description This command enables BRG processing on the specified RADIUS proxy server. Whenever an Access-Accept message is received with the attribute Alc-BRG-Id present, this will trigger the creation of a BRG. The BRG will use the **brg-profile** specified in the Access-Accept message or fall back to this BRG profile. When the specified **radius-proxy-server** has a cache enabled, no cache entries will be created for a transaction identified as BRG. A RADIUS proxy server can only be listed in one BRG profile.

This command can be executed multiple times.

The **no** form of this command removes BRG processing for the specified **radius-proxy server**.

Parameters *router-instance* — Specifies the ID of the VRF where the proxy server is located.

server-name — Specifies the name of the RADIUS proxy server.

sla-profile-string

Syntax **sla-profile-string** *string*
no sla-profile-string

Context config>subscr-mgmt>vrgw>brg>brg-profile

Description This command configures the SLA profile string that will be used as a default for SLA profile lookup. This string can be overridden during BRG or host authentication.

The **no** form of this command removes the string from the configuration.

Parameters *string* — Specifies the string to use to look up the SLA profile.

sub-profile-string

Syntax **sub-profile-string** *string*
no sub-profile-string

Context config>subscr-mgmt>vrgw>brg>brg-profile

Description This command is used as a default for subscriber profile lookup. This string can be overridden during BRG or host authentication.

The **no** form of this command removes the string from the configuration.

Parameters *string* — Specifies the string used to look up the subscriber profile.

uplink-initial-wait

Syntax **uplink-initial-wait** *seconds*

Context config>subscr-mgmt>vrgw>brg>brg-profile

Description This command specifies how long to wait for the uplink to fully establish when using a non-routed uplink such as a PPPoE client. During this initial wait time, setup of devices in the home is blocked.

After the timer expires, if an uplink was successful on only one of two IP stacks, devices continue to be configured with the successful IP stack. Control plane message for the failed IP stack are dropped.

Default uplink-initial-wait 30

Parameters *seconds* — Specifies the timeout in seconds.

Values 10 to 300

13.4.2.3 Subscriber Management Commands

pppoe-client-policy

Syntax **pppoe-client-policy** *pppoe-client-policy-name* [**create**]
no pppoe-client-policy *pppoe-client-policy-name*

Context config>subscr-mgmt

Description This command provisions a policy containing a set of parameters to be used to configure a PPPoE client.

The **no** form of this command removes the policy from the system. The policy can only be removed when it is not in use.

Parameters *pppoe-client-policy-name* — Specifies a unique name for the policy.

create — Mandatory keyword when creating a new policy.

keepalive

Syntax **keepalive** *seconds* [**hold-up-multiplier** *multiplier*]
no keepalive

Context config>subscr-mgmt>pppoe-client-policy

Description This command defines the **keepalive** interval and the number of times the **keepalive** can be missed before the session is declared down for this PPPoE client policy.

The **no** form of this command reverts to the default.

Default keepalive 30 hold-up-multiplier 3

Parameters *seconds* — Specifies the **keepalive** interval in seconds.

Values 10 to 300

multiplier — Specifies the number times **keepalive** can be missed.

Values 1 to 5

mru

Syntax **mru** *mru-bytes*
no mru

Context config>subscr-mgmt>pppoe-client-policy

Description This command defines which Maximum Receive Unit (MRU) value is signaled by the PPPoE client.

The **no** form of this command reverts to the default.

Default mru 1492

Parameters *mru-bytes* — Specifies the MRU value in octets.

Values 512 to 9154

mtu

Syntax **mtu** *mtu-bytes*
no mtu

Context config>subscr-mgmt>pppoe-client-policy

Description This command defines which Maximum Transmission Unit (MTU) is applied, by default, for packets egressing the PPP link. If a lower MRU is sent during PPP link establishment, the MRU value is used.

The **no** form of this command reverts to the default.

Default mtu 1492

Parameters *mtu-bytes* — Specifies the MTU value in octets.

Values 512 to 9154

python-policy

Syntax **python-policy** *name*
no python-policy

Context config>subscr-mgmt>pppoe-client-policy

Description This command applies a Python policy to all messages sent and received by the PPPoE client.

The **no** form of this command removes the associated Python policy from the PPPoE client.

Parameters *name* — The name of a preconfigured Python policy.

restart-backoff

Syntax **restart-backoff initial-time** *seconds* **max-time** *seconds*
no restart-backoff

Context config>subscr-mgmt>pppoe-client-policy

Description This command configures backoff timer parameters that determine how often and how long the system will attempt to restart a PPPoE client after a failure. When a client first fails, the system immediately tries to re-establish connectivity. If this attempt is also unsuccessful, the system initiates a backoff timer and waits until it expires before attempting to restart the client again, to avoid flooding the BNG. The initial duration of the backoff timer is configured with the **initial-time** parameter. With each subsequent failure, the backoff timer is doubled until the configured **max-time** is reached.

The **no** form of this command reverts to the default.

Default restart-backoff initial-time 30 max-time 600

Parameters **initial-time** *seconds* — Specifies the initial backoff time to wait before attempting a client restart.

Values 10 to 3600

max-time *seconds* — Specifies the maximum time to attempt client restarts.

Values 10 to 3600

stack

Syntax **stack** [**ipv4**] [**ipv6-slaac**]
no stack

Context config>subscr-mgmt>pppoe-client-policy

Description This command defines which NCP session is started in the PPPoE client and how addresses are retrieved within that NCP session.

Default stack ipv4

Parameters **ipv4** — Indicates that IPCP should be started and used to retrieve an IPv4 address.

ipv6-slaac — Indicates that IPv6CP should be started and that SLAAC will be used to retrieve an IPv6 prefix.

radius-accounting-policy

Syntax **radius-accounting-policy** *name* [**create**]
no radius-accounting-policy *name*

Context config>subscr-mgmt

Description This command specifies a subscriber RADIUS-based accounting policy.

The **no** form of this command removes the policy name from the configuration.

Parameters *name* — The name of the policy. The string is case-sensitive and limited to 32 ASCII 7-bit printable characters with no spaces.

create — Keyword used to create a RADIUS accounting policy instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

include-radius-attribute

Syntax [**no**] **include-radius-attribute**

Context config>subscr-mgmt>auth-plcy
config>subscr-mgmt>acct-plcy

Description This command enables the context to specify the RADIUS parameters that the system should include in RADIUS authentication-request messages.

The **no** form of this command reverts to the default values.

authentication-policy

Syntax **authentication-policy** *name* [**create**]
no authentication-policy *name*

Context config>subscr-mgmt

Description This command specifies authentication policy configured under the **aaa** context for authenticating users on the WLAN-GW ISA.

The **no** form of this command removes the policy name from the configuration.

Parameters *name* — Specifies the name of the authentication policy up to 32 characters

create — Keyword used to create a authentication policy instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

wifi-num-attached-ues

Syntax [**no**] **wifi-num-attached-ues**

Context config>subscr-mgmt>acct-plcy>include
config>subscr-mgmt>auth-plcy>include

Description This command includes the Alc-Num-Attached-UEs attribute. For session/host accounting, this is the tunnel of the associated UE. For queue-instance accounting, this attribute will not be included.

The **no** form of this command removes the attribute from inclusion.

brg-num-active-sessions

Syntax [**no**] **brg-num-active-sessions**

Context config>subscr-mgmt>acct-plcy>include

Description This command indicates the number of IPoE sessions that are currently active on the BRG to which this accounting message relates.

The **no** form of this command removes the attribute from inclusion.

subscriber-mgmt

Syntax **subscriber-mgmt**

Context config

Description This command enables the context to configure subscriber management entities. A subscriber is uniquely identified by a subscriber identification string. Each subscriber can have several DHCP sessions active at any time. Each session is referred to as a subscriber host and is identified by its IP address and MAC address.

All subscriber hosts belonging to the same subscriber are subject to the same hierarchical QoS (HQoS) processing. The HQoS processing is defined in the sub-profile (the subscriber profile). A sub-profile refers to an existing scheduler policy (configured in the **configure>qos>scheduler-policy** context) and offers the possibility to overrule the rate of individual schedulers within this policy.

Because all subscriber hosts use the same scheduler policy instance, they must all reside on the same complex.

13.4.2.4 Service Commands

brg

Syntax **brg**

Context config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>vrgw
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>vrgw
config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command enables the context to configure BRG parameters. In the **config>service>ies>sub-if>grp-if** and **config>service>vprn>sub-if>grp-if** contexts, these commands are only available in the **vlan-tag-ranges** level.

authenticated-brg-only

Syntax [**no**] **authenticated-brg-only**

Context config>service>ies>sub-if>grp-if>wlan-gw>vlan-ranges>range>vrgw>brg
config>service>vprn>sub-if>grp-if>wlan-gw>vlan-ranges>range>vrgw>brg
config>service>ies>sub-if>grp-if>brg
config>service>vprn>sub-if>grp-if>brg

Description This command indicates that only BRGs that are pre-authenticated using the RADIUS proxy are allowed in this context.

The **no** form of this command removes the restriction.

default-brg-profile

Syntax **default-brg-profile** *profile-name*
no default-brg-profile

Context config>service>vprn>sub-if>grp-if>wlan-gw>vlan-ranges>range>vrgw>brg
config>service>ies>sub-if>grp-if>wlan-gw>vlan-ranges>range>vrgw>brg
config>service>vprn>sub-if>grp-if>brg
config>service>ies>sub-if>grp-if>brg

Description This command indicates that the default BRG profile must be used for new BRGs. This profile can be overridden by RADIUS.

Parameters *profile-name* — Specifies the name of the brg-profile to be applied.

learn-ap-mac

Syntax **learn-ap-mac** [**delay-auth**]
no learn-ap-mac

Context config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw

Description This command enables the sending of ARP or ND packets on the WLAN-GW GRE tunnel for certain events. The target IP address in the ARP/ND packet is the endpoint IP address of the AP. The ARP/ND response from the AP should contain the AP MAC, which subsequently can be reported in a called-station-id message. When enabled, a message will be sent for following events:

- CPM: Mobility to an AP for which the AP-MAC is not yet known
- CPM: RS-triggered authentication on an AP for which the AP-MAC is not yet known
- ISA: Any mobility event
- ISA: Any authentication where the AP-MAC is not yet known (for example, from a RADIUS proxy cache or a DHCP circuit-id). If the optional keyword **delay-auth** is configured, then the authentication will be delayed until the ARP/ND is answered or timed out, after which the AP-MAC can be included in the authentication.

This configuration is ignored for L2-AP and L2TPv3 access.

Parameters **delay-auth** — Specifies that authentication will be delayed until the ARP/ND is answered or timed out, after which the AP-MAC can be included in the authentication.

data-triggered-ue-creation

Syntax [**no**] **data-triggered-ue-creation**

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range

Description This command enables data-triggered subscriber creation for WIFI subscribers. Data-triggered UE creation is currently only supported for UDP and TCP packets.

The **no** form of this command disables the data-triggered subscriber creation for WIFI subscribers.

subscriber-interface

Syntax **subscriber-interface** *ip-int-name* [**create**]
no subscriber-interface *ip-int-name*

Context config>service>ies
config>service>vprn

Description This command allows the operator to create special subscriber-based interfaces. It is used to contain multiple group interfaces. Multiple subnets associated with the subscriber interface can be applied to any of the contained group interfaces in any combination. The subscriber interface allows subnet sharing between group interfaces.

The **no** form of this command removes the subscriber interface.

Parameters *ip-int-name* — Specifies the interface name of a subscriber interface. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

fwd-service *service-id* — specifies the wholesale service ID.

group-interface

Syntax **group-interface** *ip-int-name* [**create**]
no group-interface *ip-int-name*

Context config>service>ies>sub-if
config>service>vprn>sub-if

Description This command creates a group interface. This interface is designed for triple-play services where multiple SAPs are part of the same subnet. A group interface may contain one or more SAPs.

The **no** form of this command removes the group interface from the subscriber interface.

Parameters *ip-int-name* — Specifies the interface name of a group interface. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

wlan-gw

Syntax **wlan-gw**

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command enables the context to configure WLAN GW parameters.

range

Syntax **range start** [*value*] **end** [*value*]
range default
no range start [*value*] **end** [*value*]

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges

config>service>ies>sub-if>grp-if>wlan-gw>ranges

Description This command creates or enters the context of specified VLAN range for configuration applicable to that range of VLANs.

Parameters **start** [0 to 4096] — Specifies the start of the VLAN range.

end [0 to 4096] — Specifies the end of the VLAN range.

default — Configures defaults for the interface.

vlan-tag-ranges

Syntax **vlan-tag-ranges**

Context config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw

Description This command enables the context to configure VLAN-to-retail-map parameters to map dot1q tags to the retail service ID. The WIFI AP inserts a dot1Q tag in the Layer 2 frame within the GRE tunnel to indicate the retail service provider for the subscriber.

xconnect

Syntax **xconnect**

Context config>router>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw>vlan-tag-ranges>range
config>service>ies>sub-if>grp-if>wlan-gw>vlan-tag-ranges>range

Description This command enables the context to configure WLAN-GW cross-connect UE-related parameters.

tunnel-source-ip

Syntax **tunnel-source-ip** *ipv6-address/prefix-length*
no tunnel-source-ip

Context config>router>wlan-gw>xconnect
config>service>vprn>wlan-gw>xconnect

Description This command configures the IPv6 address and prefix for the tunnel source.

The **no** form of this command removes IPv6 address and prefix length from the cross-connect configuration.

Parameters *ipv6-address/prefix-length* — Specifies the tunnel source IPv6 address and prefix length of the cross-connect.

Values ipv6-address/prefix: ipv6-address x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x [0..FFFF]H
d [0..255]D
(no multicast address)
prefix-length [1 to 128]

wlan-gw-group

Syntax **wlan-gw-group** *nat-group-id*
no wlan-gw-group

Context config>router>wlan-gw>cross-connectcross-connectcross-connectxconnect
config>service>vprn>wlan-gw>xconnect

Description This command configures the identifier of the WLAN Gateway ISA group that processes the cross-connect.

The no form of this command removes the NAT group IP from the cross-connect configuration.

Parameters *nat-group-id* — Specifies the identifier of the ISA group.

Values 1 to 4

accounting-policy

Syntax **accounting-policy** *isa-radius-policy-name*
no accounting-policy

Context config>service>vprn>sub-if>grp-if>wlan-gw>vlan-tag-ranges>range>xconnect
config>service>ies>sub-if>grp-if>wlan-gw>vlan-tag-ranges>range>xconnect

Description This command configures the ISA RADIUS accounting policy for the cross-connect.

The **no** form of this command removes the ISA RADIUS accounting policy from the cross-connect UE.

Parameters *isa-radius-policy-name* — Specifies the identifier of the ISA RADIUS policy name, up to 32 characters.

accounting-update-interval

Syntax **accounting-update-interval** [*interval*]
no accounting-update-interval

Context config>service>vprn>sub-if>grp-if>wlan-gw>vlan-tag-ranges>range>xconnect
config>service>ies>sub-if>grp-if>wlan-gw>vlan-tag-ranges>range>xconnect

Description This command configures the time interval between consecutive interim accounting update messages. If not configured, the system does not send interim accounting update messages.

The **no** form of this command removes the value from the cross-connect configuration.

Parameters *interval* — Specifies the time interval between consecutive interim accounting update messages in minutes.

Values 5 to 259200

mobility-acct-updates

Syntax [**no**] **mobility-acct-updates**

Context config>service>vprn>sub-if>grp-if>wlan-gw>vlan-tag-ranges>range>xconnect
config>service>ies>sub-if>grp-if>wlan-gw>vlan-tag-ranges>range>xconnect

Description This command enables the administrative state to send mobility-triggered accounting interim updates.

The **no** form of this command disables sending the mobility-triggered accounting updates.

13.4.2.5 Router Home LAN Extension Commands

vrgw

Syntax **vrgw**

Context config>router

Description This command enables the context to configure router Virtual Residential Gateway parameters.

lanext

Syntax [**no**] **lanext**

Context config>router>vrgw

Description This command enables the context to configure HLE parameters.

The **no** form of this command disables the context.

vxlan-port

Syntax **vxlan-port** [**4789** | **8472**]
no vxlan-port

Context config>router>vrgw>lanext

Description This command specifies the destination UDP port for both ingress and egress VXLAN packets for HLE services.

The **no** form of this command reverts to the default.

Default vxlan-port 4789

Parameters **4789 | 8472** — Specifies the destination UDP port.

Values 4789, 8472

vxlan-vtep-range

Syntax **vxlan-vtep-range start** *ip-address* | *ipv6-address* **end** *ip-address* | *ipv6-address*
no vxlan-vtep-range

Context config>router>vrgw>lanext

Description This command specifies the address range to be used as the local VXLAN VTEP on the ISA for HLE services. The system allocates one address for each ISA in the WLAN GW group out of the specified range. The allocated address appears as /32 or /128 routes in the global routing table with the route type **nat**.

The **no** form of this command removes the values from the configuration.

Parameters **start** — Specifies the start of the VXLAN VTEP range.

end — Specifies the end of the VXLAN VTEP range.

ip-address **|** *ipv6-address* — Specifies the range of VXLAN VTEP addresses.

Values <ip-address|ipv6-address : ipv4-address - a.b.c.d
ipv6-address - x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x - [0..FFFF]H
d - [0..255]D

wlan-gw-group

Syntax **wlan-gw-group** *wlan-gw-group-id*
no wlan-gw-group

Context config>router>vrgw>lanext
config>router>wlan-gw>xconnect

Description This command specifies the WLAN GW group that is used for HLE services.

The **no** form of this command removes the group from the configuration.

Parameters *wlan-gw-group-id* — Specifies the WLAN GW group ID.

Values 1 to 4

13.4.2.6 Subscriber Management Commands

radius-accounting-policy

Syntax **radius-accounting-policy** *name* [**create**]
no radius-accounting-policy *name*

Context config>subscr-mgmt

Description This command specifies a subscriber RADIUS-based accounting policy.

The **no** form of this command removes the policy name from the configuration.

Parameters *name* — The name of the policy. The string is case-sensitive and limited to 32 ASCII 7-bit printable characters with no spaces.

create — Keyword used to create a RADIUS accounting policy instance. The **create** keyword requirement can be enabled/disabled in the **environment>create** context.

include-radius-attribute

Syntax [**no**] **include-radius-attribute**

Context config>subscr-mgmt>acct-plcy

Description This command enables the context to specify the RADIUS parameters that the system should include in RADIUS authentication-request messages.

The **no** form of this command disables any RADIUS parameters enabled in this context.

lanext-bridge-id

Syntax [**no**] **lanext-bridge-id**

Context config>subscr-mgmt>acct-plcy>include

Description This command enables the system to include the HLE service's bridge ID (Alc-Bridge-Id) in RADIUS accounting packets.

The **no** form of this command excludes the HLE service's bridge ID (Alc-Bridge-Id) from RADIUS accounting packets.

lanext-device-type

Syntax [**no**] **lanext-device-type**

Context config>subscr-mgmt>acct-plcy>include

Description This command enables the system to include the HLE host's device type (Alc-HLE-Device-Type) in RADIUS accounting packets.

The **no** form of this command excludes the HLE host's device type (Alc-HLE-Device-Type) from RADIUS accounting packets.

lanext-route-distinguisher

Syntax [**no**] **lanext-route-distinguisher**

Context config>subscr-mgmt>acct-plcy>include

Description This command enables the system to include the HLE service's EVPN route distinguisher (Alc-RD) in RADIUS accounting packets.

The **no** form of this command excludes the HLE service's EVPN route distinguisher (Alc-RD) from RADIUS accounting packets.

lanext-route-target

Syntax [**no**] **lanext-route-target**

Context config>subscr-mgmt>acct-plcy>include

Description This command enables the system to include the HLE service's EVPN route target (Alc-RT) in RADIUS accounting packets.

The **no** form of this command excludes the HLE service's EVPN route target (Alc-RT) from RADIUS accounting packets.

lanext-vni

Syntax [**no**] **lanext-vni**

Context config>subscr-mgmt>acct-plcy>include

Description This command enables the system to include the HLE service's EVPN VXLAN VNI (Alc-Vxlan-VNI) in RADIUS accounting packets.

The **no** form of this command excludes the HLE service's EVPN VXLAN VNI (Alc-Vxlan-VNI) from RADIUS accounting packets.

vrgw

Syntax **vrgw**

Context config>subscr-mgmt

Description This command enables the context to configure router Virtual Residential Gateway (vRGW). parameters.

lanext

Syntax [**no**] **lanext**

Context config>subscr-mgmt>vrgw

Description This command enables the context to configure subscriber management vRGW home HLE parameters.

The **no** form of this command disables the context.

router-target-as-number

Syntax **router-target-as-number** *as-number*
no router-target-as-number

Context config>subscr-mgmt>vrgw>lanext

Description This command specifies the AS number for the HLE service. It is used to derive the route target (RT) and route distinguisher (RD) for the HLE EVPN service only when the RADIUS server does not return a specific route target or route distinguisher.

The derived RT is in the "target:<configured-router-target-as-number>:<returned Alc-Bridge-Id>" format.

The derived RD is in the “<configured-router-target-as-number>:<returned Alc-Bridge-Id>” format.

The **no** form of this command removes the AS number from the configuration.

Parameters *as-number* — Specifies the AS number of the HLE service.

Values 1 to 65535

13.4.2.7 Service Commands

subscriber-interface

Syntax **subscriber-interface** *ip-int-name* [**create**]
no subscriber-interface *ip-int-name*

Context config>service>ies
config>service>vprn

Description This command allows the operator to create special subscriber-based interfaces. It is used to contain multiple group interfaces. Multiple subnets associated with the subscriber interface can be applied to any of the contained group interfaces in any combination. The subscriber interface allows subnet sharing between group interfaces.

The **no** form of this command removes the subscriber interface.

Parameters *ip-int-name* — Specifies the interface name of a subscriber interface. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

group-interface

Syntax **group-interface** *ip-int-name* [**create**]
no group-interface *ip-int-name*

Context config>service>ies>sub-if
config>service>vprn>sub-if

Description This command creates a group interface. This interface is designed for triple-play services where multiple SAPs are part of the same subnet. A group interface can contain one or more SAPs.

The **no** form of this command removes the group interface from the subscriber interface.

Parameters *ip-int-name* — Specifies the interface name of a group interface. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

wlan-gw

Syntax **wlan-gw**

Context config>service>ies>sub-if>grp-if
config>service>vprn>sub-if>grp-if

Description This command enables the context to configure WLAN GW parameters.

max-lanext-bd

Syntax **max-lanext-bd** [*value*]
no max-lanext-bd

Context config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw

Description This command specifies the maximum number of HLE BDs for this group interface.

The **no** form of this command disables HLE for the group interface.

Parameters *value* — Specifies the maximum number of Bridged Domains for this interface.

Values 1 to 131071

vlan-tag-ranges

Syntax **vlan-tag-ranges**

Context config>service>ies>sub-if>grp-if>wlan-gw
config>service>vprn>sub-if>grp-if>wlan-gw

Description This command enables the context to configure IEEE 802.1q VLAN tag ranges.

range

Syntax **range start** [*value*] **end** [*value*]
range default
no range start [*value*] **end** [*value*]

Context config>service>vprn>sub-if>grp-if>wlan-gw>ranges

config>service>ies>sub-if>grp-if>wlan-gw>ranges

Description This command creates a VLAN range or enters the context of the specified VLAN ranges for configuration applicable to that range of VLANs.

Parameters **start** — Specifies the start of the VLAN range.

Values 0 to 4096

Default 200

end — Specifies the end of VLAN range.

Values 0 to 4096

Default 400

default — Specifies to use defaults for the interface.

vrgw

Syntax **vrgw**

Context config>service>ies>sub-if>grp-if>wlan-gw>ranges>range
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range

Description This command enables the context to configure Virtual Residential Gateway parameters.

lanext

Syntax [**no**] **lanext**

Context config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>vrgw
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>vrgw

Description This command enables the context to configure HLE parameters.

The **no** form of this command disables the vRGW parameters enabled in this context.

access

Syntax [**no**] **access**

Context config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext

Description This command enables the context to configure the access side of HLE for the VLAN range.

The **no** form of this command disables the vRGW parameters enabled in this context.

max-mac

Syntax **max-mac** [*value*]
no max-mac

Context config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext>access
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext>access
config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext>network
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext>network

Description This command specifies the maximum number of allowed MAC addresses on the access side of HLE.

The **no** form of this command reverts to the default.

Default max-mac 20

Parameters *value* — Specifies the maximum number of MAC entries in bridged domains.

Values 1 to 128

multi-access

Syntax [**no**] **multi-access**

Context config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext>access
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext>access

Description This command enables access from multiple APs into a per-tenant BD and the associated vRGW (BRG) instance.

The **no** form of this command disables access from multiple APs and limits access from a single AP into per tenant bridge domain (BD) and the associated vRGW (BRG) instance.

assistive-address-resolution

Syntax [**no**] **assistive-address-resolution**

Context config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext

Description This command enables assistive address resolution (AAR) for HLE services.

bd-mac-prefix

Syntax **bd-mac-prefix** *mac-prefix*
no bd-mac-prefix

Context config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext

Description This command specifies the prefix of the HLE BD MAC address.

The **no** form of this command removes the MAC prefix from the configuration.

mac-translation

Syntax [**no**] **mac-translation**

Context config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext

Description This command enables MAC address translation for HLE services.

The **no** form of this command disables MAC address translation for HLE services.

network

Syntax [**no**] **network**

Context config>service>ies>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext
config>service>vprn>sub-if>grp-if>wlan-gw>ranges>range>vrgw>lanext

Description This command enables the context to configure network side attributes.

The **no** form of this command resets the network parameters to the default values.

13.4.2.8 Show Commands

vrgw

Syntax **vrgw**

Context show>subscr-mgmt

Description This command displays Virtual Residential Gateway information.

brg

Syntax **brg**

Context show>subscr-mgmt>vrgw

Description This command enables the context to display operational information for connected BRGs and related hosts.

brg-hosts

Syntax **brg-hosts** [**mac** *ieee-address*] [**ip** *ip-address*]

Context show>subscr-mgmt>vrgw>brg

Description This command displays operational information for connected BRGs and related hosts.

Parameters *ieee-address* — Displays detailed information for each prefix.

ip-address — Displays only the prefixes associated with this subscriber interface.

Output The following output is an example of subscriber management BRG host information.

Sample Output

```
Node# show subscriber-mgmt brg brg-hosts
===============================================================================
Bridged Residential Gateway hosts
===============================================================================
Identifier                  : 00:00:00:00:00:05
MAC address                 : 00:02:00:00:00:19
IP address                  : 21.0.0.20
Service                     : 4 (VPRN)
Allocation type             : dynamic
Home-aware pool             : 00:00:00:00:00:05
DHCP lease                  : true
Remaining lease time        : 482
Lease start time            : 2016/01/18 17:44:44
Identifier                  : 00:00:00:00:00:05
MAC address                 : 00:02:00:00:00:19
IP address                  : 3ffe:0:0:5::
Service                     : 4 (VPRN)
Identifier                  : 00:00:00:00:00:05
MAC address                 : 00:02:00:00:00:1a
IP address                  : 21.0.0.21
Service                     : 4 (VPRN)
Allocation type             : dynamic
Home-aware pool             : 00:00:00:00:00:05
DHCP lease                  : true
Remaining lease time        : 486
Lease start time            : 2016/01/18 17:44:48
Identifier                  : 00:00:00:00:00:05
MAC address                 : 00:02:00:00:00:1a
IP address                  : 3ffe:0:0:5::
Service                     : 4 (VPRN)
Identifier                  : 00:00:00:00:00:05
MAC address                 : 00:02:00:00:00:1b
IP address                  : 21.0.0.22
Service                     : 4 (VPRN)
```

```
Allocation type              : dynamic
Home-aware pool              : 00:00:00:00:00:05
DHCP lease                   : true
Remaining lease time         : 490
Lease start time             : 2016/01/18 17:44:52
-------------------------------------------------------------------------------
No. of BRG hosts: 5
===============================================================================
```

gateway

Syntax **gateway brg-id** *brg-ident*

Context show>subscr-mgmt>vrgw>brg

Description This command displays all operational data related to the BRG.

Parameters *brg-ident* — Specifies the BRG identifier, up to 32 characters.

Output The following output is an example of subscriber management BRG gateway information.

Sample Output

```
Node# show subscriber-mgmt brg gateway brg-id "00:00:00:00:00:05"
===============================================================================
Bridged Residential Gateways
===============================================================================
Identifier                   : 00:00:00:00:00:05
SLAAC prefix                 : 3ffe:0:0:5::/64
Subnet                       : 21.0.0.1/24
Subnet start address         : 21.0.0.20
Subnet end address           : 21.0.0.40
DMZ address                  : (Not Specified)
DNS 1 v4                     : (Not Specified)
DNS 1 v6                     : (Not Specified)
DNS 2 v4                     : (Not Specified)
DNS 2 v6                     : (Not Specified)
NBNS 1                       : (Not Specified)
NBNS 2                       : (Not Specified)
DHCP lease time              : 600
DHCP stream destination      : (Not Specified)
IPv4 portal URL              : (Not Specified)
IPv6 portal URL              : (Not Specified)
BRG profile                  : brg
Subscriber profile           : subprof_mig_4
SLA profile                  : (Not Specified)
UPnP policy override         : (Not Specified)
DMZ address in use           : no
Proxy authenticated          : yes
Ingress IPv4 filter override : N/A
Egress IPv4 filter override  : N/A
Ingress IPv6 filter override : N/A
Egress IPv6 filter override  : N/A
No QoS overrides found.
No Filter rules received.
```

```
===============================================================================
```

bindings

Syntax **bindings** [**mac** *ieee-address*]

Context show>subscr-mgmt>vrgw>brg>gateway

Description This command lists all the DHCP bindings of the local pool associated with the specified BRG. If the MAC attribute is specified, only the binding for that MAC address is displayed.

Parameters *ieee-address* — specifies the MAC address identifying a specific binding.

Output The following output is an example of BRG gateway binding information.

Sample Output

```
Node# show subscriber-mgmt brg gateway brg-id "00:00:00:00:00:05" bindings
===============================================================================
Bridged Residential Gateway home-aware pool address bindings
===============================================================================
Home-aware pool              : 00:00:00:00:00:05
-------------------------------------------------------------------------------
MAC address                  : 00:02:00:00:00:19
IP address                   : 21.0.0.20
Allocation type              : dynamic
DHCP lease                   : true
Remaining lease time         : 489
Lease start time             : 2016/01/18 17:44:44
MAC address                  : 00:02:00:00:00:1a
IP address                   : 21.0.0.21
Allocation type              : dynamic
DHCP lease                   : true
Remaining lease time         : 493
Lease start time             : 2016/01/18 17:44:48
MAC address                  : 00:02:00:00:00:1b
IP address                   : 21.0.0.22
Allocation type              : dynamic
DHCP lease                   : true
Remaining lease time         : 497
Lease start time             : 2016/01/18 17:44:52
-------------------------------------------------------------------------------
No. of bindings: 3
===============================================================================
```

host

Syntax **host mac** *ieee-address* **ip** *ip-address*

Context show>subscr-mgmt>vrgw>brg>gateway

Description This command lists pool information about a specific host in the BRG.

Parameters *ieee-address* — Specifies the MAC address identifying a specific host.

ip-address — Displays only the prefixes associated with this host.

hosts

Syntax **hosts**

Context show>subscr-mgmt>vrgw>brg>gateway

Description This command lists all the hosts connected to the BRG.

Output The following output is an example of BRG gateway host information.

Sample Output

```
Node# show subscriber-mgmt brg gateway brg-id "00:00:00:00:00:05" hosts
===============================================================================
Bridged Residential Gateway hosts
===============================================================================
Identifier                  : 00:00:00:00:00:05
MAC address                 : 00:02:00:00:00:19
IP address                  : 21.0.0.20
Service                     : 4 (VPRN)
Allocation type             : dynamic
Home-aware pool             : 00:00:00:00:00:05
DHCP lease                  : true
Remaining lease time        : 388
Lease start time            : 2016/01/18 17:44:44
Identifier                  : 00:00:00:00:00:05
MAC address                 : 00:02:00:00:00:19
IP address                  : 3ffe:0:0:5::
Service                     : 4 (VPRN)
Identifier                  : 00:00:00:00:00:05
MAC address                 : 00:02:00:00:00:1a
IP address                  : 21.0.0.21
Service                     : 4 (VPRN)
Allocation type             : dynamic
Home-aware pool             : 00:00:00:00:00:05
DHCP lease                  : true
Remaining lease time        : 392
Lease start time            : 2016/01/18 17:44:48
Identifier                  : 00:00:00:00:00:05
MAC address                 : 00:02:00:00:00:1a
IP address                  : 3ffe:0:0:5::
Service                     : 4 (VPRN)
Identifier                  : 00:00:00:00:00:05
MAC address                 : 00:02:00:00:00:1b
IP address                  : 21.0.0.22
Service                     : 4 (VPRN)
Allocation type             : dynamic
Home-aware pool             : 00:00:00:00:00:05
DHCP lease                  : true
Remaining lease time        : 396
Lease start time            : 2016/01/18 17:44:52
```

```
-------------------------------------------------------------------------------
No. of BRG hosts: 5
===============================================================================
```

standby-ip-addresses

Syntax **standby-ip-addresses**

Context show>subscr-mgmt>vrgw>brg

Description This command shows associated home-aware pool standby IP addresses.

gateways

Syntax **gateways**
gateways brg-profile *profile-name*

Context show>subscr-mgmt>vrgw>brg

Description This command lists operational gateway information for all BRGs active on the system. This list can optionally be reduced by specifying a BRG profile.

Parameters *profile-name* — Specifies the name of the BRG profile to filter.

Output The following output is an example of BRG gateway information.

Sample Output

```
Node# show subscriber-mgmt brg gateways
===============================================================================
Bridged Residential Gateways
===============================================================================
Identifier                  : 00:00:00:00:00:05
SLAAC prefix                : 3ffe:0:0:5::/64
Subnet                      : 21.0.0.1/24
Subnet start address        : 21.0.0.20
Subnet end address          : 21.0.0.40
DMZ address                 : (Not Specified)
DNS 1 v4                    : (Not Specified)
DNS 1 v6                    : (Not Specified)
DNS 2 v4                    : (Not Specified)
DNS 2 v6                    : (Not Specified)
NBNS 1                      : (Not Specified)
NBNS 2                      : (Not Specified)
DHCP lease time             : 600
DHCP stream destination     : (Not Specified)
IPv4 portal URL             : (Not Specified)
IPv6 portal URL             : (Not Specified)
BRG profile                 : brg
Subscriber profile          : subprof_mig_4
SLA profile                 : (Not Specified)
UPnP policy override        : (Not Specified)
```

```
DMZ address in use           : no
Proxy authenticated          : yes
Ingress IPv4 filter override : N/A
Egress IPv4 filter override  : N/A
Ingress IPv6 filter override : N/A
Egress IPv6 filter override  : N/A
No QoS overrides found.
No Filter rules received.
-------------------------------------------------------------------------------
No. of gateways: 1
===============================================================================
```

lanext

Syntax **lanext**

Context show>subscr-mgmt>vrgw

Description This command displays Wireless LAN Gateway Home LAN Extension information.

pppoe-client

Syntax **pppoe-client** [*brg-id* [*32 chars max*]]
pppoe-client [*brg-id* [*32 chars max*]] [**detail**]
pppoe-client [*brg-id* [*32 chars max*]] [**statistics**]

Context show>subscr-mgmt>vrgw>brg

Description This command displays information of PPPoE clients started by the vRGW.

Parameters *brg-id* — Specifies the string identifying the BRG.

detail — Displays detailed PPPoE client information.

statistics — Displays PPPoE client statistics.

Output The following output is an example of PPPoE client information.

Sample Output

```
Node# show subscriber-mgmt vrgw brg pppoe-client
===============================================================
PPPoE Clients
===============================================================
Id
  Svc-Id     MAC               PPPoE-Sid Up-Time
    IP-Address      IPv6-Prefix
---------------------------------------------------------------
00:0c:29:00:00:10
  20         00:0c:29:00:00:10 1         0d 00:00:05
    10.255.1.16     2001:db8:101::/64
---------------------------------------------------------------
```

```
No. of PPPoE Clients: 1
==============================================================
Node# show subscriber-mgmt vrgw brg pppoe-client detail
==============================================================================
PPPoE Clients detail
==============================================================================
Id                  : 00:0c:29:00:00:10
Service Id          : 20
MAC Address         : 00:0c:29:00:00:10
Up Time             : 0d 00:00:27
Policy              : clnt_plcy
PPPoE Session Id    : 1
Service Name        : (Not Specified)
AC MAC              : 08:06:ff:00:00:02
AC Name             : BNG
LCP State           : opened
IPCP State          : opened
IPv6CP State        : opened
Local Magic Num     : 201366211
Remote Magic Num    : 374038276
Remote MRU          : 1492
PPP User-Name       : 00:0c:29:00:00:10
PPP Auth-Protocol   : none
IP Address          : 10.255.1.16
Primary DNS         : 1.1.1.1
Secondary DNS       : 1.1.2.2
Primary NBNS        : 2.2.1.1
Secondary NBNS      : 2.2.2.2
IPv6 Intf Id        : 02:0c:29:ff:fe:00:00:10
IPv6 Prefix         : 2001:db8:101::/64
Prefix Timeout      : 86375
Primary IPv6 DNS    : 2001:db8:dddd:1::1
Secondary IPv6 DNS  : 2001:db8:dddd:2::1
------------------------------------------------------------------------------
No. of PPPoE Clients: 1
==============================================================================
Node# show subscriber-mgmt vrgw brg pppoe-client statistics
================================================
PPPoE Clients statistics
================================================
Id             : 00:0c:29:00:00:10
Packet Type          Received       Transmitted
PADI                 -              1
PADO                 1              -
PADR                 -              1
PADS                 1              -
PADT                 0              0
Session              9              9
LCP Echo Request
(since restart)      0              21
LCP Echo Reply
(since restart)      21             0
RS                   -              1
RA                   1              -
DHCP6 Info           -              1
DHCP6 Reply          1              -
------------------------------------------------
Rx Dropped    : 0
------------------------------------------------
```

```
No. of PPPoE Clients: 1
===============================================
```

brg-profile

Syntax **brg-profile**
brg-profile *name*
brg-profile *name* **associations**

Context show>subscr-mgmt

Description This command displays BRG profile information.

Parameters *name* — Specifies an existing BRG profile name.

associations — Displays information associated with the specified BRG profile name.

bd

Syntax **bd** [*bridge-id*]

Context show>subscr-mgmt>vrgw
show>subscr-mgmt>vrgw>lanext

Description This command displays information for the specified HLE bridge. If no bridge is specified, then all domains are listed.

Parameters *bridge-id* — Specifies the BD ID.

Values 1 to 4294967295

Output The following output is an example of HLE bridge information.

Sample Output

```
show subscriber-mgmt vrgw lanext bd 100
===============================================================================
BRIDGE DOMAIN(S)
===============================================================================
BRIDGE-ID                    : 100
VNI                          : 1
ROUTE TARGET                 : TARGET:100:100
ROUTE DISTINGUISHER          : 100:100
WLANGW GRPID                 : 1
ISA MEMBERID                 : 1
WLANGW BD VLANTAG            : N/A
WLANGW BD SERVICE            : 500
WLANGW BD INTERFACE          : G1
WLANGW BD MAC TRANSLATION    : N/A
WLANGW BD MAC                : AA:BB:CC:00:00:01
WLANGW BD ASSIST.ADDR.RES    : DISABLED
```

```
WLANGW BD NETW MAXMAC        : 20
WLANGW BD NETW STATE         : ENABLED
WLANGW BD ACCS MAXMAC        : 20
-------------------------------------------------------------------------------
NO. OF BRIDGE DOMAINS: 1
===============================================================================
```

13.4.2.9 Debug Commands

brg-id

Syntax [**no**] **brg-id** *brg-ident*

Context debug>subscr-mgmt>vrgw>brg>pppoe-client>brg-id

Description This command enables debugging of PPPoE client messages linked to a BRG.

Parameters *brg-ident* — The string identifying the BRG.

detail-level

Syntax **detail-level** {**low** | **medium** | **high**}

Context debug>subscr-mgmt>vrgw>brg>pppoe-client>brg-id

Description This command specifies the amount of detail present in debugging the specified PPPoE client.

Default detail-level high

Parameters **low** — Specifies a low level of detail during debugging.

medium — Specifies a medium level of detail during debugging.

high — Specifies a high level of detail during debugging.

direction

Syntax **direction** {**ingress** | **egress** | **both**}

Context debug>subscr-mgmt>vrgw>brg>pppoe-client>brg-id

Description This command specifies if debugging should only include ingress, egress or all messages.

Default direction both

Parameters **ingress** — Specifies that debugging only includes ingress messages.

egress — Specifies that debugging only includes egress messages.

both — Specifies that debugging includes both ingress and egress messages.

discovery

Syntax **discovery** [**padi**] [**pado**] [**padr**] [**pads**] [**padt**]

Context debug>subscr-mgmt>vrgw>brg>pppoe-client>brg-id

Description This command, limits debugging only to the specified messages in the discovery phase.

Parameters **padi** — Limits debugging only to padi messages.

pado — Limits debugging only to pado messages.

padr — Limits debugging only to padr messages.

pads — Limits debugging only to pads messages.

padt — Limits debugging only to padt messages.

dropped-only

Syntax [**no**] **dropped-only**

Context debug>subscr-mgmt>vrgw>brg>pppoe-client>brg-id

Description This command specifies that only packets that are dropped by the vRGW will be shown in debugging.

Default dropped-only

ppp

Syntax **ppp** [**lcp**] [**pap**] [**chap**] [**ipcp**] [**ipv6cp**] [**ipv6**]

Context debug>subscr-mgmt>vrgw>brg>pppoe-client>brg-id

Description This command specifies which messages in PPP setup are tracked by debugging. If no messages are specified, they are all tracked. LCP Echo Request and Echo Response are never shown during debugging.

Parameters **lcp** — Tracks lcp messages during debugging.

pap — Tracks pap messages during debugging.

chap — Tracks chap messages during debugging.

ipcp — Tracks ipcp messages during debugging.

ipv6cp — Tracks ipv6cp messages during debugging.

ipv6 — Tracks ipv6 messages during debugging.

13.4.2.10 Clear Commands

vrgw

Syntax **vrgw**

Context clear>subscr-mgmt

Description This command clears vRGW data.

brg

Syntax **brg**

Context clear>subscr-mgmt>vrgw

Description This command clears BRG data.

gateway

Syntax **gateway brg-id** *brg-ident*
gateway brg-id *brg-ident* **idle-bindings** [**binding** *ieee-address*]
gateway all-gateways
gateway brg-id *brg-ident* **all-hosts**
gateway brg-id *brg-ident* **host** *ieee-address*

Context clear>subscr-mgmt>vrgw>brg

Description This command clears BRG gateway data.

Parameters *brg-ident* — Clears the specified BRG and all related hosts from the system.

idle-bindings — Clears a specific idle binding or all idle binding from the BRG. An idle binding is a DHCP lease that has not been expired but for which the host has been removed due to idle timeout.

ieee-address — Specifies the MAC address of the binding to remove.

all-gateways — Clears all BRGs and related hosts from the system.

all-hosts — Clears all hosts belonging to a specified BRG.

ieee-address — Clears a specific host belonging to a specified BRG.

13.4.2.11 Tools Commands

restart

Syntax **restart brg-id** *brg-ident*

Context tools>perform>subscr-mgmt>vrgw>brg>pppoe-client

Description This command forcibly restarts the PPPoE client linked to the BRG instance. The restart does not remove BRG or related session state, however, forwarding is impacted.

Parameters *brg-ident* — The string identifying the BRG.

lanext

Syntax **lanext**

Context tools>dump>wlan-gw

Description This command enables the context to dump tools information for WGLAN HLE parameter.

bd

Syntax **bd** *bridge-id*

Context tools>dump>wlan-gw>lanext
tools>perform>wlan-gw>lanext

Description This command enables the context to enter the specified HLE BD to perform tools commands.

Parameters *bridge-id* — Specifies the HLE BD ID.

Values 1 to 4294967295

arp-table

Syntax **arp-table** [**ip** *ipv4-address*] [**class** *class*] [**next-index** *index*]

Context tools>dump>wlan-gw>lanext>bd

Description This command dumps the specified ARP table entries for the specified HLE BD.

Parameters *ipv4-address* — Specifies the IP address of the ARP entry to be displayed.

class — Specifies the source on which to filter the ARP table entries.

Values access, network, remote

index — Specifies the next set of results if there are more results than the current output. The additional results are shown at the bottom of the current output.

Values 0 to 4294967295

Output The following output is an example of ARP table information.

Sample Output

```
tools>dump>wlan-gw>lanext>bd# arp-table
===============================================================================
MATCHED 1 ARP ENTRY ON SLOT #2 MDA #1
===============================================================================
IP ADDRESS       MAC ADDRESS        CLASS    TIMESTAMP
-------------------------------------------------------------------------------
192.168.0.1      AA:BB:CC:00:00:01  ACCESS   05/30/2017 23:49:11
===============================================================================
```

flood-table

Syntax **flood-table** [**class** *class*] [**detail**]

Context tools>dump>wlan-gw>lanext>bd

Description This command dumps the specified flood table entries for the specified HLE BD.

Parameters *class* — Specifies the source on which to filter the flood table entries.

Values access, network, remote

detail — Dumps detailed information for the flood table.

Output The following output is an example of flood table information.

Sample Output

```
tools>dump>wlan-gw>lanext>bd# flood-table detail
===============================================================================
MATCHED 2 FLOOD ENTRIES ON SLOT #2 MDA #1
===============================================================================
TUNNEL MDA      : 2/1                   VLAN            : 100
TUNNEL PORT     : 2/1/NAT-IN-IP         TUNNEL ENCAP    : 2049.1
TUNNEL SRC IP   : 11.3.3.100
TUNNEL DST IP   : 4.4.4.4
TUNNEL TYPE     : GRE                   TUNNEL CLASS    : ACCESS
GRE RAW HEADER  : 1483014144
-------------------------------------------------------------------------------
TUNNEL MDA      : 2/1                   VLAN            : N/A
TUNNEL PORT     : 2/1/NAT-OUT-IP        TUNNEL ENCAP    : 2081.3
```

```
TUNNEL SRC IP   : 2.2.2.2
TUNNEL DST IP   : 44.44.44.1
TUNNEL TYPE     : VXLAN                 TUNNEL CLASS    : NETWORK
VXLAN VNI       : 102
VXLAN PORT      : 4789
-------------------------------------------------------------------------------
===============================================================================
```

mac-table

Syntax **mac-table** [**mac** *ieee-address*] [**class** *class*] [**next-index** *index*] [**detail**]

Context tools>dump>wlan-gw>lanext>bd

Description This command dumps the specified MAC table entries for the specified HLE BD.

Parameters *ieee-address* — Specifies the MAC address of the MAC table entry.

class — Specifies the source on which to filter the MAC table entries.

Values access, network, remote

index — Specify the **next-index** *index* parameter to display the next set of results If there are more results than the current output. The additional results are indicated at the bottom of the current output.

Values 0 to 4294967295

Output The following output is an example of MAC table information.

Sample Output

```
tools>dump>wlan-gw>lanext>bd# mac-table
===============================================================================
MATCHED 2 MAC ENTRIES ON SLOT #2 MDA #1
===============================================================================
UE-MAC            VLAN   BRIDGE-ID  DESCRIPTION            L2-SVC     L2-VLAN
  TUNNEL  CLASS   TYPE   ROUTER     SRC-IP
  ANCHOR                            DST-IP
-------------------------------------------------------------------------------
AA:2B:9E:11:D1:6F 100    100        ESM-USER               N/A        N/A
  2/1     ACCESS  GRE    500        11.3.3.100
  2/1/NAT-OUT-IP:2145.1             4.4.4.4
-------------------------------------------------------------------------------
02:DE:01:00:01:11 N/A    100        ESM-USER               N/A        N/A
  2/1     NETWORK VXLAN  500        2.2.2.2
  2/1/NAT-OUT-IP:2145.1             44.44.44.1
-------------------------------------------------------------------------------
===============================================================================
```

neighbor-table

Syntax **neighbor-table** [**ip6** *ipv6-address*] [**class** *class*] [**next-index** *index*]

Context tools>dump>wlan-gw>lanext>bd

Description This command dumps neighbor table entries for the specified HLE BD.

Parameters *ipv6-address* — Specifies the IPv6 address of the neighbor entry to be displayed.

class — Specifies the source on which to filter the neighbor table entries.

Values access, network, remote

index — Specify the **next-index** *index* parameter to display the next set of results If there are more results than the current output. The additional results are indicated at the bottom of the current output.

Values 0 to 4294967295

lanext

Syntax **lanext**

Context tools>perform>wlan-gw

Description This command enables the context to clear tools information for WGLAN HLE parameter.

clear-arp

Syntax **clear-arp** [**ip** *ipv4-address*]

Context tools>perform>wlan-gw>lanext>bd

Description This command clears all ARP table entries for the specified HLE BD or a specific entry if an IPv4 address is specified.

Parameters *ipv4-address* — Specifies the IPv4 address of the ARP table entry to be cleared.

clear-mac

Syntax **clear-mac** [**mac** *ieee-address*]

Context tools>perform>wlan-gw>lanext>bd

Description This command clears all MAC table entries for the specified HLE BD or a specific entry if a MAC address is specified.

Parameters *ieee-address* — Specifies the MAC address of the MAC table entry to be cleared.

clear-neighbor

Syntax **clear-neighbor** [**ip6** *ipv6-address*]

Context tools>perform>wlan-gw>lanext>bd

Description This command clears all neighbor table entries for the specified HLE BD or a specific entry if an IPv6 address is specified.

Parameters *ipv6-address* — Specifies the IPv6 address of the neighbor table entry to be cleared.

14 Service Chaining for ESM Hosts with L2-Aware NAT

This section describes service chaining for ESM hosts with L2-Aware NAT.

14.1 Steering to Service Chains for ESM hosts with L2-Aware NAT

This feature provides the steering of traffic flows for ESM hosts with L2-aware NAT (typically used in vRGW and WLAN-GW), to service-functions (SF) in a data-center, reachable over an IP underlay network with a tunnel. The supported tunnel encapsulation is VxLAN. The steering function on the gateway (vRGW or WLAN-GW) is configured through a PBR action in an ISA filter (also referred to as VAS filter) associated with an L2-aware NAT host. The PBR action specifies the IP address of the SF, the EVPN service instance through which the SF is reached, and optionally, an ESI. A controller in the data-center, and the subscriber edge (a gateway such as vRGW, WLAN-GW, or BNG) acting as a service-function-classifier (SFC) participate in BGP-EVPN to exchange reachability information for the SF in the DC, and NAT pools on the gateway. The gateway resolves the PBR target, in other words, SF IP@ in the EVPN service configured in the VAS filter via BGP EVPN routes received from the controller in the DC. The network virtualization edge (NVE) in the DC, such as a host stack running the SF in a VM can act as a bridge or a router. Nuage VRS or VSG is an example of an NVE.

The ISA filter used for steering can also be configured with an optional action to insert a network services header (NSH) in steered traffic, as described in RFC 8300, *Network Service Header (NSH)*.

A group of NAT ISAs that provide per-host or per-flow steering functions for L2-aware NAT hosts on the gateway are configurable under the base router. The steered traffic from the gateway to the VAS is VxLAN encapsulated on these ISAs. A range of IP addresses used for local VxLAN VTEPs on ISAs is configured (and routable) under the base router.

14.1.1 Terminology

Controller — A network element in the DC that peers with the subscriber edge to exchange BGP EVPN routes. The controller can advertise type-1 (ESI), type-2 (MAC advertisement, with an optional IP address), and type-5 (an IP prefix route containing an SF's IP address or subnet) routes for each SF. The controller also learns EVPN type-2 and type-5 routes from the subscriber edge for reachability of NAT pools on the ISAs. The controller can program the NVE with these routes.

EVPN import-mode — A configurable attribute of the EVPN service on the subscriber edge that controls the type of EVPN route information that is imported or tracked from the EVPN peer and exported or advertised to the EVPN peer. The possible values are **bridged**, **routed**, or **none**. For import mode **bridged**, the subscriber edge imports EVPN type-2 routes containing the SF's MAC and IP addresses and the associated VxLAN VTEP to reach the SF. With import mode **routed**, the EVPN type-5 route containing the SF's IP address or subnet and the type-2 route containing the NVE's IP and MAC address are imported or tracked. The NVE, in this case, routes traffic to or from the SF. Also, with **import-mode routed**, the subscriber edge advertises type-2 and type-5 routes for NAT prefixes (outside pools). If no import-mode is configured, then only type-2 and type-5 routes for NAT prefixes are advertised to the peer (no routes are imported). See EVPN Route Updates and Tracking.

NSH — Network Services Header. Extra encapsulation carried in the steered traffic over VxLAN tunnels that contains information about the packet handling through the chain of value added service functions in the DC. In addition to the information about the service chain that the steered traffic traverses, it also contains optional meta-data (for example, the subscriber-id). More details on NSH usage and format are defined in VAS Filters on the ISA.

NVE — Network virtualization edge. A network element in the DC (for example, TOR or virtual switch or router such as the Nuage VRS) that provides connectivity to the SFs (VM or appliance). The NVE can perform bridging or routing functions for traffic to or from the SF. EVPN service instances (including R-VPLS tied to a VPRN) can be configured on the NVE, depending on the configured forwarding mode (bridging or routing to connected SFs).

Service-chaining service — An EVPN control plane instance on the subscriber edge that is used to learn and track EVPN routes for reachability to the SFs in the DC. It is also used to advertise EVPN routes for NAT pools on the subscriber edge to the controller in the DC.

SF — Service function. The value-added service (for example, a parental control service) running in a VM or on an appliance in the Data Center (DC). The SF can be part of a chain of SFs providing opt-in value added services.

SFC — Service function classifier. The component on the vRGW/WLAN-GW/BNG (subscriber edge) that matches flows belonging to the L2-aware NAT host that has opted-in for value-added services, and performs steering (PBR) functions to the first service in the service chain that is applicable to the host. SFC supports optional NSH insertion in the steered traffic. Flow matching and steering is specified in VAS filters that are applied upstream and downstream flows on the ISA. SFC also tracks EVPN routes (type-1, type-2, and type-5) to resolve the SF's IP address, and optionally, ESI in an EVPN service, all of which is specified in the steering action in the VAS filters.

14.2 VAS Filters on the ISA

A VAS filter can be associated with individual hosts of an L2-aware NAT subscriber either at creation time or via CoA. The name of the VAS filter associated with the host is provided by the AAA server. This VAS filter is applied on each packet that creates a new flow. This results in a pair of actions which are then used for all subsequent upstream and downstream packets in that flow. These actions allow forwarding different kinds of traffic to different service chains in both upstream and downstream directions.

14.2.1 Matching

A VAS filter contains a set of ordered entries. Each entry contains a single match criteria and either an upstream action, a downstream action, or both. Matching occurs independently for both directions, only taking into account the entries for which an action is configured in that direction. Matching stops at the first matching entry. Actions from different entries for either direction can result.

The VAS filter entries can match on:

- a foreign IP or subnet (for example, a destination IP and subnet for upstream traffic, and source IP and subnet for downstream traffic match)
- a foreign port (for example, a destination port for upstream and a source port for downstream traffic)
- a protocol

14.2.2 Forwarding

The resulting forwarding action can be one of the following:

- **action forward** — The packet is forwarded as usual to its original destination.
- **action forward sf-ip** *sf-ip* [**esi** *esi*] **service-id** *svc-id* — The packet is tunneled toward the **sf-ip** over EVPN/VxLAN overlay.

  If the **sf-ip** cannot be resolved (see EVPN Route Updates and Tracking), the fail-action is used. This can be either forward or drop. The default is forward. Optionally, an NSH header can be inserted.

14.2.3 NSH Insertion

NSH insertion is an optional action of the VAS filter when forwarding to the service chain. An additional header is inserted between the tunneling Ethernet header and the IP payload, as defined in RFC 8300, *Network Service Header (NSH)*. VxLAN-GPE is not supported.

IP => UDP (port 4789) => VxLAN => Ethernet =>NSH => IP

The NSH header is populated with following values:

- TTL in the base header is set to 63.
- The **service-path-id** (24-bit number) and **service-index** (8-bit number) are filled in from the filter entry's action.
- The MD-Type is set to 1 and the 16-byte metadata is filled in. The source of this metadata is discussed below. If no metadata is provided, this field is zero.

The metadata can be specified in the filter entry's action. The metadata can either be a 16-byte opaque data hex string (zero-padded if it is smaller than 16 bytes), or it can be derived from the subscriber string (in the Alc-Subsc-Id-Str VSA). In the latter case, the subscriber string is truncated after the first 16 bytes, and thus, these first 16 bytes should be unique.

Alternatively, the opaque data string can be provided by AAA. This source has precedence over the filter entry's action.

14.2.4 Configuration

Modifications to an existing VAS filter are applicable to all hosts using that VAS filter.

- Adding filter entries is possible at all times and the new entry is taken into account on the next filter application at flow creation. Existing flows are not re-evaluated against the modified filter.

- Removing filter entries is only possible when the filter entry is shut down.
- Modifying the match criteria on a filter entry is only possible when the filter entry is shut down.
- Adding or removing a filter entry action is only possible when the filter entry is shut down.
- Modifying an existing filter entry's action is possible at all times and takes effect immediately for all flows currently matched against that filter entry action.

When a filter entry is shut down, all flows matched against any of the filter entry actions is killed. The **clear subscriber-mgmt isa-service-chaining vas-filter** *filter-name* command can be executed to clear all flows of all hosts associated with the filter. This command can be used if all flows needing reevaluation against a modified set of matching criteria.

A maximum of 512 VAS filters with up to 64 entries per filter are supported.

14.3 EVPN Route Updates and Tracking

14.3.1 NVE Bridging to SF

In this scenario, the NVE acts as a bridge for upstream traffic. The service chain (EVPN) service on subscriber edge is configured with import-mode bridged. In this mode, the SF IP address in the VAS filter action is resolved to obtain SF's MAC address and VxLAN VTEP/VNI to reach it. If the filter action only specifies SF IP address and no ESI, then EVPN route type-2 (MAC route with MAC and IP address of SF) is used to resolve SF's IP address. If a route is found, then the SF MAC address and VxLAN VTEP/VNI from the type-2 route are used. If ESI is specified in the filter action, then the ESI is resolved using EVPN type-1 route. In this case, the MAC address is taken from route type-2, and VxLAN and VNI are taken from route type-1. The resolution is downloaded to the ISA. The upstream packet from the host is VxLAN encapsulated. The destination MAC in the inner Ethernet frame is the SF's MAC address. The source MAC is ISA's MAC address.

```
config>sub-mgmt>service-chaining
     service-chain 500 import-mode bridged create
         bgp
            route-distinguisher 65001:500
            route-target import 65000:500
         no shut
```

A single separate EVPN service (a backbone EVPN) can be configured between the controller and subscriber edge. This is used to advertise type-5 routes for NAT pools on the subscriber edge.

```
config>sub-mgmt>service-chaining
       service-chain 600 create
          bgp
             route-distinguisher 65001:600
             route-target export 65000:600
          evpn
             vxlan vni 10
             nat-group 1
             gw-address-range start 120.1.1.10 end 120.1.1.4
             ip-advertise-routes ipv4 nat
                 outside-svc 400 outside-pool pool-1
                 outside-svc 400 outside-pool pool-2
```

A Layer 3 domain (overlay VPRN) must be configured on the NVE. The subscriber edge is connected to this Layer 3 domain on the NVE via an EVPN (the backbone EVPN or R-VPLS on NVE). The R-VPLS interface on the NVE (into this VPRN) must share the same subnet with the subscriber edge. NVE receives the type-5 routes in the backbone EVPN (R-VPLS), and adds the received type-5 prefixes to the FIB in the VPRN that the R-VPLS is connected to. The gateway IP address in the EVPN

type-5 routes sent by the subscriber edge must be on the same subnet as the R-VPLS interface on NVE. A range of gateway IP addresses in this subnet are configured under the EVPN service on the subscriber edge, such that each individual ISA gets a gateway IP address to use for exported type-5 routes. NVE acts as a router for downstream traffic from the SF that is destined to NAT outside IP address of the subscriber (Figure 200).

Figure 200 **NVE Bridging Traffic to the SF**

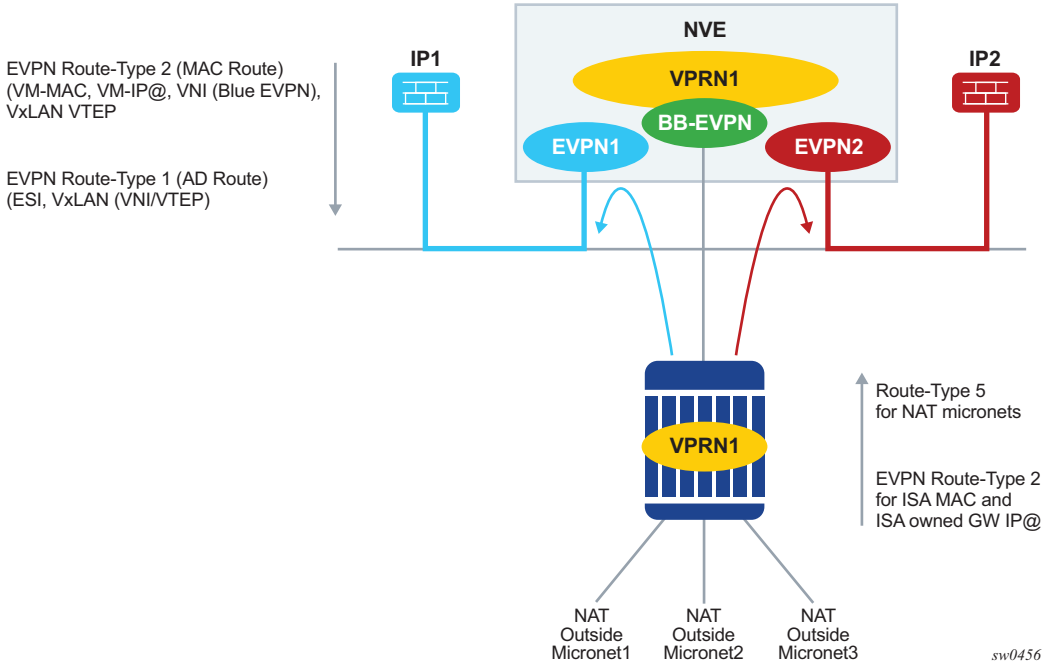

14.3.2 NVE Routing to SF

In this situation, the NVE acts as a router for both upstream (subscriber edge to SF) and downstream traffic (SF to the subscriber edge). The SF must be configured with NVE as the default router.

The service-chain (EVPN) service on subscriber edge is configured with import-mode **routed**. NAT pools are advertised in EVPN type-5 routes. The ISA's gateway IP address is sent in a type-5 route (Figure 201).

Figure 201 **EVPN Route Type-5**

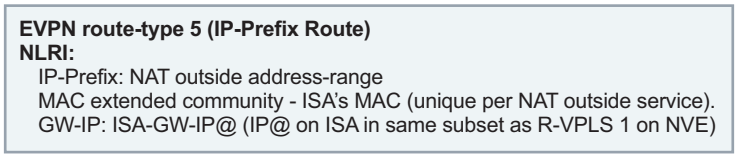

A type-2 router advertising ISA's MAC address and gateway IP address is generated (Figure 202).

Figure 202 **EVPN Route Type-2**

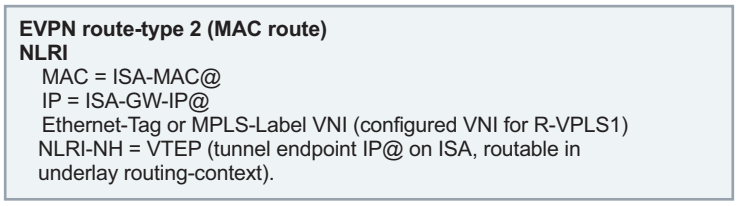

This EVPN service (called a backbone EVPN) is configured on NVE as an R-VPLS that ties to a VPRN as described in NVE Routing to SF. NVE adds the received type-5 prefixes to the FIB in VPRN.

The controller in the data-center also generates type-5 routes carrying IP address or subnet for the SF. The MAC address in a type-5 route must be the R-VPLS interface MAC or NVE's system MAC address (since the NVE is routing traffic). NVE should also generate an EVPN route type-2 to advertise the MAC of the R-VPLS interface (or single MAC address of the NVE), and optionally the IP address of the R-VPLS interface. This is shown in Figure 203.

On subscriber edge, SF's IP address in the filter action is resolved in the configured service in the filter action via a best match IP lookup against EVPN route type-5. If the resolved route type-5 has non zero GW-IP, then a recursive lookup (exact match) is done in the service. If it is resolved by EVPN route type-2, then the nexthop MAC (DA MAC that is used in inner Ethernet header) and the VxLAN VNI/VTEP are taken from the route type-2. If GW-IP in route type-5 is zero, then the dest MAC and nexthop (VTEP) is take from the route type-5.

Figure 203 **NVE Routing Traffic To and From the SF**

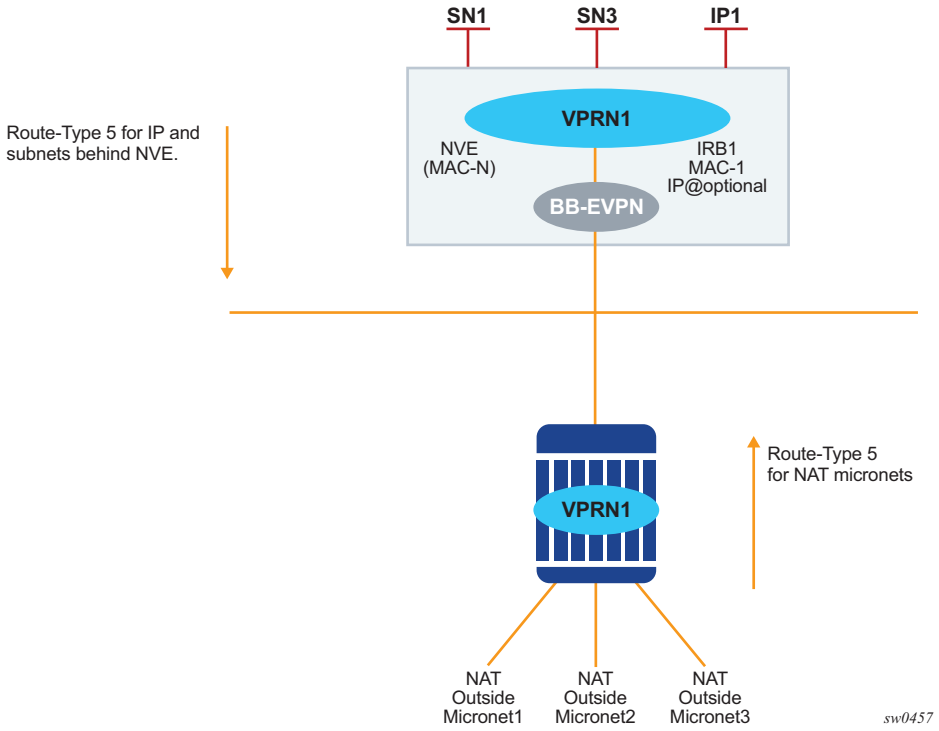

14.4 Data Path on the Subscriber Edge

14.4.1 Upstream Traffic (Access to Network)

In this scenario, L2-aware NAT is performed. Traffic is subjected to VAS filtering associated with the host. Based on the match entry, the output of the action contains an SF IP address, EVPN service instance, optional ESI, and optional NSH parameters (**service-path-id**, **service-index**, and optional **metadata**). SF IP address and optional ESI are resolved in the indicated EVPN service as per the configured import-mode of the EVPN service (described in the previous sections). The result of resolution is SF MAC address, VxLAN VTEP and VNI. The upstream packet is encapsulated as shown below:

- With an optional NSH insertion, the encapsulation used is VxLAN, where the Ethernet header following the VxLAN carries Ether-Type NSH, as shown below. VxLAN-GPE is not supported.

  IP => UDP (port 4789) => VxLAN => Ethernet => NSH => IP

  - Outer IP source address: Local VTEP (ISA's local IP address)
  - Outer IP destination address: Remote VTEP (from EVPN route, as described in previous sections)
  - Destination MAC in Ethernet header: SF MAC address or NVE's MAC address (depending on bridging or routing on NVE).
  - Source MAC in Ethernet header: MAC address of ISA (generated based on configured MAC prefix)
  - VNI in VxLAN: VNI from EVPN route resolution (as defined in previous sections)
  - Inner IP (payload) source address: NAT outside IP address
  - Inner IP (payload) destination address: original destination on the Internet
  - Without NSH the encapsulation is standard VxLAN, for example, IP => UDP => VxLAN => Ethernet => IP. The parameters in outer IP, inner Ethernet, and IP are as defined about (in other words, the same as for the encapsulation with NSH).

14.4.2 Downstream Traffic — From Network

The destination IP address in the downstream traffic is the subscriber's NAT outside IP address. The traffic is revived on the right NAT ISA. NAT flow is looked up. Associated VAS filter action is executed. The action can be steering (with or without NSH), with steering parameters, such as an SF IP address, EVPN service instance, (optional) ESI, and (optional) NSH parameters (**service-path-id**, **service-index**, optional **meta-data**). SF IP address and optional ESI are resolved in the indicated EVPN service as per the configured import-mode of the EVPN service (described in previous sections). The result of resolution is SF MAC address, VxLAN VTEP and VNI. The downstream packet steered to the SF is encapsulated similarly to upstream traffic described in NVE Bridging to SF.

14.5 Data Path on NVE

Upstream traffic from the subscriber edge is either bridged or routed by the NVE. The EVPN service is located based on the VNI. The destination MAC address in the Ethernet frame is either the SF's MAC address, in which case the traffic is bridged to the SF behind the NVE, or the destination MAC matches the local R-VPLS interface MAC address or single MAC address of the NVE. Then, the packet is routed (based on FIB lookup on the NVE) to the right SF.

Downstream traffic (network to subscriber edge to vas) after being processed by the VAS chain contains the NAT outside IP address in the IP destination address. SFs are configured with NVE as the default router. The packet is received on the NVE from the SF. A lookup is performed in the FIB on NVE (which contains IP prefixes populated by EVPN route type-5 updates). The next-hop IP address is the IP address of the ISA (GW-IP field in route type-5). An ARP and L2FDB lookup respectively provides the destination MAC address, and VxLAN VTEP and VNI. Traffic is VxLAN encapsulated and forwarded in the underlay toward the subscriber edge.

14.6 Service Chaining for ESM Hosts with L2-Aware NAT

14.6.1 Command Hierarchies

• ISA Service Chaining Commands
• Show Commands

14.6.1.1 ISA Service Chaining Commands

```
configure
    — router
        — [no] isa-service-chaining
            — nat-group nat-group-id [create]
            — no nat-group nat-group-id
            — vxlan-vtep-range start [ip-address | ipv6-address] end [ip-address | ipv6-address]
            — no vxlan-vtep-range
    — subscriber-mgmt
        — isa-service-chaining
            — evpn service-id [import-mode import-mode] [create]
            — no evpn service-id
                — [no] bgp
                    — route-distinguisher auto-rd
                    — route-distinguisher rd
                    — no route-distinguisher
                    — route-target export ext-community import ext-community
                    — no route-target
                — description description-string
                — no description
                — [no] export
                    — gw-address-range start start end end
                    — no gw-address-range
                    — ip-advertise-routes
                        — pool pool-name router router-instance [create]
                        — no pool pool-name router router-instance
                        — [no] shutdown
                    — vxlan vni vni
                    — no vxlan
                — [no] shutdown
            — mac-prefix mac-prefix
            — no mac-prefix
            — vas-filter name [create]
            — no vas-filter name
                — description description-string
```

```
— no description
— entry id [create]
— no entry id
    — action direction [create]
    — no action direction
        — fail-action fail-action
        — no fail-action
        — forward
        — forward sf-ip ip-address | ipv6-address svc service-id
            [esi esi]
        — no forward
        — [no] insert-nsh
            — meta-data
                — [no] insert-subscriber-id
                — opaque-data hex-string
                — no opaque-data
            — svc-path path-id svc-index service-index
            — no svc-path
    — description description-string
    — no description
    — match
        — foreign-ip ip-address[lmask]
        — no foreign-ip
        — foreign-port port
        — no foreign-port
        — protocol protocol-id
        — no protocol
    — [no] shutdown
```

14.6.1.2 Show Commands

```
show
— subscriber-mgmt
    — isa-service-chaining
        — mappings svc service-id [sf-ip ip-address | ipv6-address] [esi esi] [state map-
            state]
        — vas-filter name [entry [entry]]
        — vas-filter
— router
    — isa-service-chaining
```

14.6.2 Command Descriptions

14.6.2.1 Generic Commands

description

Syntax **description** *description-string*
no description

Context config>subscr-mgmt>isa-svc-chain>evpn
config>subscr-mgmt>isa-svc-chain>vas-filter

Description This command creates a text description stored in the configuration file for a configuration context. This associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes the description string from the context.

Parameters *description-string* — Specifies a text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

shutdown

Syntax [**no**] **shutdown**

Context config>subscr-mgmt>isa-svc-chain>evpn
config>subscr-mgmt>isa-svc-chain>vas-filter

Description This command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command.

The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they can be deleted.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system-generated configuration files.

The **no** form of this command puts an entity into the administratively enabled state.

14.6.2.2 Service Chaining Commands

isa-service-chaining

Syntax [**no**] **isa-service-chaining**

Context config>router

Description This command enables the context to configure ISA service chaining parameters.

The **no** form of the command disables ISA service chaining parameters.

nat-group

Syntax **nat-group** *nat-group-id* [**create**]
no nat-group nat-group-id

Context config>router>isa-svc-chain

Description This command allows service chaining to be enabled for subscribers whose NAT flows are established on the set of ISAs in the specified NAT group.

The **no** form of this command removes the NAT group from the configuration.

Parameters *nat-group-id* — Specifies the NAT group identifier.

Values 1 to 4

create — Keyword used to create the NAT group instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

vxlan-vtep-range

Syntax **vxlan-vtep-range start** [*ip-address* | *ipv6-address*] **end** [*ip-address* | *ipv6-address*]
no vxlan-vtep-range

Context config>router>isa-svc-chain

Description This command specifies the address range to be used as the local VXLAN VTEP on the ISA for service chaining. The system allocates one address for each ISA in the NAT group out of the specified range. The allocated address appears as /32 or /128 routes in the global routing table with the route type **nat**.

The **no** form of this command removes the IP or IPv6 addresses from the configuration.

Parameters **start** — Specifies the start address of the VXLAN VTEP range.

end — Specifies the end address of the VXLAN VTEP range.

ip-address — Specifies an IPv4 address start and end range.

Values ipv4-address - a.b.c.d

ipv6-address — Specifies an IPv6 address start and end range

ipv6-address: x:x:x:x:x:x:x:x (eight 16-bit pieces)
x:x:x:x:x:x:d.d.d.d
x - [0..FFFF]H
d - [0..255]D

isa-service-chaining

Syntax **isa-service-chaining**

Context config>subscr-mgmt

Description This command enables the context to configure ISA-based service chaining for subscribers with L2-Aware NAT.

evpn

Syntax **evpn** *service-id* [**import-mode** *import-mode*] [**create**]
no evpn *service-id*

Context config>subscr-mgmt>isa-svc-chain

Description This command configures the import mode for the service chaining EVPN service. The **import-mode** controls the EPVN route types that are imported by the EVPN system.

The **no** form of this command removes the configuration parameters.

Parameters *service-id* — Specifies the service ID of the EVPN.

Values 1 to 2147483647

import-mode — Specifies the import mode of the EVPN.

Values bridged — The specified EVPN instance imports EVPN route type-2 and type-1 from the peer.

routed — The specified EVPN instance imports EVPN type-1, type-2 and type-5 routes from the peer. Also, the EVPN instance can be configured to export EVPN type-5 routes for NAT pools to the peer.

none — The specified EVPN instance does not import any EVPN routes from the peer but can be configured with NAT pools that are exported to the peer in EVPN type-5 routes.

create — Keyword used to create the EVPN service instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

bgp

Syntax [**no**] **bgp**

Context config>subscr-mgmt>isa-svc-chain>evpn

Description This command enables the context to configure EVPN BGP-specific information.

The **no** form of this command reverts to the default.

route-distinguisher

Syntax **route-distinguisher auto-rd**
route-distinguisher *rd*
no route-distinguisher

Context config>subscr-mgmt>isa-svc-chain>evpn>bgp

Description This command configures the Route Distinguisher (RD) field that is signaled in NLRI in EVPN routes.

The **no** form of this command reverts to the default.

Parameters **auto-rd** — Specifies that the system automatically generates an RD.

rd — Specifies the RD.

Values rd: *ip-addr:comm-val* | *2byte-asnumber:ext-comm-val* | *4byte-asnumber:comm-val*
ip-addr: a.b.c.d
comm-val: [0 to 65535]
2byte-asnumber: [1 to 65535]
ext-comm-val: [0 to 4294967295]
4byte-asnumber: [1 to 4294967295]

route-target

Syntax **route-target export** *ext-community* **import** *ext-community*
no route-target

Context config>subscr-mgmt>isa-svc-chain>evpn>bgp

Description This command configures route target attributes to be signaled in EVPN routes used for service chaining.

The **no** form of this command removes the parameters from the configuration.

Parameters **export** — Specifies the route target to be used by BGP in this EVPN service when exporting EVPN routes.

import — Specifies the route target to be used by BGP in this EVPN service when importing EVPN routes.

ext-community —

Values rd: *ip-addr:comm-val* | *2byte-asnumber:ext-comm-val* | *4byte-asnumber:comm-val*

ip-addr: a.b.c.d

comm-val: [0 to 65535]

2byte-asnumber: [1 to 65535]

ext-comm-val: [0 to 4294967295]

4byte-asnumber: [1 to 4294967295]

export

Syntax [**no**] **export**

Context config>subscr-mgmt>isa-svc-chain>evpn

Description This command enables the context to configure information related to the export of EVPN BGP routes related to service chaining.

The **no** form of this command disables exporting EVPN BGP routes related to service chaining

gw-address-range

Syntax **gw-address-range start** *start* **end** *end*
no gw-address-range

Context config>subscr-mgmt>isa-svc-chain>evpn>export

Description This command specifies the address range to be used for the gateway IP address field in EVPN type-5 routes that are advertised for configured NAT pools, to the peer for service-chaining. The system allocates one address for each ISA in the NAT group out of the specified range.

The **no** form of this command removes the values from the configuration..

Parameters *start* — Specifies the starting gateway address range (V4) for this EVPN service.

Values ipv4-address: a.b.c.d

end — Specifies the ending gateway address range (V4) for this EVPN service.

Values ipv4-address: a.b.c.d

ip-advertise-routes

Syntax **ip-advertise-routes**

Context config>subscr-mgmt>isa-svc-chain>evpn>export

Description This command enables the context to configure EVPN routes to be advertised to a BGP EVPN peer participating in service chaining.

pool

Syntax **pool** *pool-name* **router** *router-instance* [**create**]
no pool *pool-name* **router** *router-instance*

Context config>subscr-mgmt>isa-svc-chain>evpn>ip-routes

Description This command configures NAT pools that are advertised in EVPN type 5 routes to the peer participating in service chaining.

The **no** form of this command removes the parameters from the configuration.

Parameters *pool-name* — Specifies the name of the NAT pool up, to 32 characters.

router-instance — Specifies the router instance belonging to the pool.

Values *router-name* | *vprn-svc-id*
router-name: B ase, management, *cpm-vr-name,* vpls-management
Default - Base
vprn-svc-id: [1 to 2147483647]
cpm-vr-name: [up to 32 characters]
service-name: [up to 64 characters]

create — Keyword used to create a pool instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

vxlan

Syntax **vxlan vni** *vni*
no vxlan

Context config>subscr-mgmt>isa-svc-chain>evpn>export

Description This command configures a VxLAN VNI that is sent in EVPN routes advertised to the service chaining.

The **no** form of this command removes the VNI from the configuration.

Parameters *vni* — Specifies the VNI of the VXLAN created by the EVPN service.

Values 1 to 16777215

mac-prefix

Syntax **mac-prefix** *mac-prefix*
no mac-prefix

Context config>subscr-mgmt>isa-svc-chain

Description This command configures the unique MAC prefix per ISA and per outside service for all NAT group configured for service-chaining.

The **no** form of this command removes the MAC prefix from the configuration.

Parameters *mac-prefix* — Specifies the MAC prefix, up to eight characters, including separators.

Values format AA:BB:CC

vas-filter

Syntax **vas-filter** *name* [**create**]
no vas-filter *name*

Context config>subscr-mgmt>isa-svc-chain

Description This command configures a Value Added Service filter.

The **no** form of this command removes the VAS filter name from the configuration.

Default This command configures a value added service (VAS) filter that can be associated to an L2-aware NAT host, and is matched on the NAT ISA to select flows for a host that needs to be steered to remote value-added services.

Parameters *name* — Specifies a VAS filter name, up to 32 characters.

create — Keyword used to create the VAS filter instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

entry

Syntax **entry** *id* [**create**]
no entry *id*

Context config>subscr-mgmt>isa-svc-chain>vas-filter

Description This command configures an entry in the VAS filter.

The **no** form of this command removes the entry ID from the configuration.

Parameters *id* — Specified an entry in the VAS filter.

Values 0 to 4294967295

create — Keyword used to create the entry ID instance. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

action

Syntax **action** *direction* [**create**]
no action *direction*

Context config>subscr-mgmt>isa-svc-chain>vas-filter>entry

Description This command enables the context to configure an action to be performed for traffic that matches a configured match criteria in the filter entry. The action can be configured as being applicable to upstream traffic, downstream traffic, or both.

The **no** form of this command removes the direction from the configuration.

Parameters *direction* — Specifies the direction for the action in a VAS filter entry.

Values upstream, downstream

create — Keyword used to create the action's direction. The **create** keyword requirement can be enabled or disabled in the **environment>create** context.

fail-action

Syntax **fail-action** *fail-action*
no fail-action

Context config>subscr-mgmt>isa-svc-chain>vas-filter>entry>action

Description This command configures the fail action when a packet matches with a VAS filter entry in a specific direction, but no mapping exists for the specified SF-IP or ESI in the specified EVPN service.

The **no** form of this command removes the fail action from the configuration.

Parameters *fail-action* — Specifies the fail action.

Values drop, forward

forward

Syntax **forward**

forward sf-ip *ip-address* | *ipv6-address* **svc** *service-id* [**esi** *esi*]
no forward

Context config>subscr-mgmt>isa-svc-chain>vas-filter>entry>action

Description This command configures the forward action.

The **no** form of this command removes the parameters from the configuration..

Parameters *ip-address* — Specifies forwarding the SF IPv4 address for the action in a VAS filter entry.

ipv6-address — Specifies forwarding the SF IPv6 address for the action in a VAS filter entry.

service-id — Specifies the service ID.

Values 1 to 2147483647

esi — Specifies the ESI for the action in a VAS filter entry.

Values 10-byte Ethernet Segment Identifier:
00-11-22-33-44-55-66-77-88-99 with any of these separators ('-',':',' ')

insert-nsh

Syntax [**no**] **insert-nsh**

Context config>subscr-mgmt>isa-svc-chain>vas-filter>entry>action

Description This command enables the context to configure NSH parameters in the steered traffic.

The **no** form of this command removes insert NSA parameters from the configuration.

meta-data

Syntax **meta-data**

Context config>subscr-mgmt>isa-svc-chain>vas-filter>entry>action>insert-nsh

Description This command enables the context to configure opaque metadata to be inserted in NSH in the steered traffic if the forward action indicates NSH insertion.

insert-subscriber-id

Syntax [**no**] **insert-subscriber-id**

Context config>isa-svc-chain>vas-filter>entry>action>insert-nsh>meta-data

Description This command specifies that the metadata to be inserted in NSH (with MD-Type set to 1) must contain a subscriber identifier that is derived from the subscriber string that comes from the AAA server (in Alc-Subsc-Id-Str VSA). The subscriber string is truncated after the first 16 bytes, and therefore, the first 16 bytes should be unique. The **insert-subscriber-id** and **insert-subscriber-id** commands are mutually exclusive.

The **no** form of this command reverts to the default.

opaque-data

Syntax **opaque-data** *hex-string*
no opaque-data

Context config>isa-svc-chain>vas-filter>entry>action>insert-nsh>meta-data

Description This command specifies 16-byte opaque data HEX string to be inserted in NSH meta-data (with MD-Type set to 1). The opaque data can also be provided (overridden) by AAA server. AAA server has precedence over static configuration. The **opaque-data** and **insert-subscriber-id** commands are mutually exclusive

The **no** form of this command removes the HEX string from the configuration.

Parameters *hex-string* — Specifies the HEX string up to 32 characters.

svc-path

Syntax **svc-path** *path-id* **svc-index** *service-index*
no svc-path

Context config>isa-svc-chain>vas-filter>entry>action>insert-nsh

Description This command configures the service path identifier and service index to be inserted in NSH in the steered traffic if the forward action indicates NSH insertion.

The **no** form of this command removes the parameters from the configuration.

Parameters *path-id* — Specifies the 24-bit path ID in the base part of NSH.

Values 0 to 16777215

service-index — Specifies the 8-bit service index inserted in the base part of NSH.

Values 0 to 255

match

Syntax **match**

Context config>subscr-mgmt>isa-svc-chain>vas-filter>entry

Description This command enables the context to configure the match criterion for a VAS filter entry.

foreign-ip

Syntax **foreign-ip** *ip-address*[*/mask*]
no foreign-ip

Context config>subscr-mgmt>isa-svc-chain>vas-filter>entry>match

Description This command configures the foreign IP address or subnet in the match criterium for this entry. The foreign IP or subnet implies a matching destination IP for upstream traffic and a source IP for downstream traffic.

The **no** form of this command removes the IP address or subnet from the match criterium in the entry

Parameters *ip-address/mask* — Specifies the IPv4 address and mask.

Values ip-address a.b.c.d
mask 0 to 32

foreign-port

Syntax **foreign-port** *port*
no foreign-port

Context config>subscr-mgmt>isa-svc-chain>vas-filter>entry>match

Description This command configures the foreign TCP/UDP port to match in this entry of the VAS filter.

The **no** form of this command

Parameters *port* — Specifies the foreign IP port to match.

Values 1 to 65535

protocol

Syntax **protocol** *protocol-id*
no protocol

Context config>subscr-mgmt>isa-svc-chain>vas-filter>entry>match

Description This command configures the protocol ID to be matched in this entry of the VAS filter.

The **no** form of this command removes the protocol ID from the match criterium in the entry.

Parameters *protocol-id* — Specifies the protocol to match.

Values protocol-id: *protocol-number* | *protocol-name*
protocol-number: 1, 6, 17]D
[0x1,0x6,0x11]H [0b1,0b110,0b10001]B
protocol-name: none, icmp, tcp, udp
* udp/tcp wildcard

14.6.2.3 Show Commands

mappings

Syntax **mappings svc** *service-id* [**sf-ip** *ip-address* | *ipv6-address*] [**esi** *esi*] [**state** *map-state*]

Context show>subscr-mgmt>isa-svc-chain

Description This command shows reachability details for the SF in the specified EVPN instance. This information is gathered from the EVPN routes learned from the peer.

Parameters *service-id* — Displays information about the specified service ID.

Values 1 to 2147483647

ip-address — Displays information about the specified IPv4 address up to 64 characters.

ipv6-address — Displays information about the specified IPv6 address up to 64 characters.

esi — Displays information about the specified the ESI value.

Values 10-byte Ethernet Segment Identifier
00-11-22-33-44-55-66-77-88-99 with
any of these separators ('-',':',' ')

map-state — Displays map state status.

Values resolved, unresolved

Output The following is an example of service chaining mapping information.

Sample Output

```
show subscriber-mgmt isa-service-chaining mappings svc 111
===============================================================================
ISA Service Chaining - Mappings for EVPN 111
===============================================================================
IP-mappings
           MAC                Vni       Vtep
-------------------------------------------------------------------------------
160.0.0.7
```

```
           00:00:01:00:00:01  1         5.20.1.12
170.0.0.7
           00:00:02:00:00:01  1         5.20.1.12
-------------------------------------------------------------------------------
ESI-mappings
                              Vni       Vtep
-------------------------------------------------------------------------------
24:23:2c:2e:2f:36:2e:4d:1b:4c
                              1         5.20.1.12
-------------------------------------------------------------------------------
No. of mappings: 3
```

Table 167 **Service Chaining Mappings Field Descriptions**

Label	Description
IP-mappings	The IP address
MAC	The MAC address of the IP mappings
Vni	The VxLAN network identifier
Vtep	The VxLAN tunnel endpoint
ESI-mappings	The Ethernet Segment Identifier MAC address
No. of mappings	The total number of mappings for the EVPN ID.

vas-filter

Syntax **vas-filter** *name* [**entry** [*entry*]]
vas-filter

Context show>subscr-mgmt>isa-svc-chain

Description This command displays VAS filter information.

Parameters *name* — Displays information about the specified name of the VAS filter.

entry — Displays information about the specified VAS filter entry

Values 0 to 4294967295

Output The following is an example of VAS filter information.

Sample Output

```
show subscriber-mgmt isa-service-chaining vas-filter "foo"
===============================================================
ISA Service Chaining - VAS Filter "vas-filter"
===============================================================
Description            : (Not Specified)
-------------------------------------------------------------------------------
```

```
Entry                       : 1
Admin State            : in-service
Description            : (Not Specified)
Match
    Protocol             :   udp
    Foreign IP          :  12.1.1.1/24
    Foreign Port       :  2000
Action upstream
    Fail Action        : forward
    Forwarding
        Forward IP     : 40.1.0.1
        Service           : 11
        ESI                 : N/A
     Network Service Header
        Service Path   : 32
        Service Index  : 2
        Insert Sub-Id  : disabled
        Meta Data      : 0xaabbccdd
      Action downstream
          Fail Action        : forward
          Forwarding
              Forward IP     : 30.1.0.1
             Service           : 11
              ESI                : N/A
         Network Service Header
             Service Path   : 32
             Service Index  : 2
             Insert Sub-Id  : disabled
             Meta Data      : 0xaabbccdd
```

Table 168 **Service Chaining VAS Filter Field Descriptions**

Label	Description
Description	The text string describing the VAS filter
Entry	The ID of the entry in the VAS filter
Admin State	The administrative state of the entry in the VAS filter
Match	Protocol — The IP protocol to match in this entry of the VAS filter Foreign IP — The foreign IP address to match in this entry of the VAS filter Foreign Port — The foreign port to match in this entry of the VAS filter
Action	upstream — Indicates the upstream direction for the action in the VAS filter entry downstream — Indicates the downstream direction for the action in the VAS filter entry

Table 168 **Service Chaining VAS Filter Field Descriptions (Continued)**

Label	Description
Fail Action	Forwarding IP — The forwarding Service Function IP address for the action in a VAS filter entry Service — The forwarding service for the action in a VAS filter entry ESA — The ESI for the action in a VAS filter entry
Network Service Header	Service Path — The path ID used to construct an NSH (Network Service Header) for the action in a VAS filter entry
	Service Index — The service ID used to construct an NSH for the action in a VAS filter entry
	Insert Sub-ID disabled — The subscriber ID is not inserted with meta data to construct an NSH for the action in a VAS filter entry enabled — The subscriber ID is inserted with meta data to construct an NSH for the action in a VAS filter entry
	Meta Data — The opaque metadata used to construct an NSH (Network Service Header) for the action in a VAS filter entry

isa-service-chaining

Syntax **isa-service-chaining**

Context show>router

Description This command displays information related to service chaining (for ESM hosts with L2-aware NAT) at the router level.

Output The following is an example of ISA service chaining information.

Sample Output

```
show router isa-service-chaining
===============================================================
ISA Service Chaining - Router "Base"
===============================================================
VXLAN Tep range        : 1.2.3.4 - 1.2.3.8
NAT Groups                  : 2
===============================================================
```

Table 169

Label	Description
VXLAN Tep range	The VXLAN Tep Information (VTEP) range
NAT Groups	The NAT ISA group used in a virtual router for Service Chaining.

15 Dynamic Data Services

15.1 Introduction to Dynamic Data Services

Dynamic Data Services enables a zero-touch, single-ended provisioning model for business services. Two deployment models are available:

- **With control channel**: triggered by the authentication of a single or dual-stack PPPoE or IPoE session as a business CPE control channel, parameters are passed in a RADIUS Access-Accept or CoA message to set up a Layer 2 or Layer 3 data service.
- **Without control channel**: RADIUS or local authentication of a Dynamic Data Service data trigger provides the required parameters to set up a Layer 2 or Layer 3 data service.

Dynamic Data Services support includes:

- Epipe VLL services
- Epipe VLL services with dynamic MS-PWs (FEC 129) or spoke SDP
- EVPN based Epipe service
- VPLS services
- VPLS services with BGP-AD pseudowire
- VPLS service with spoke SDP or mesh SDP
- EVPN based VPLS service
- IES and VPRN services
- MVPN services
- Routing policy, filter, operational groups and OAM-PM provisioning

The full list of supported configuration commands can be displayed with the **tools dump service dynamic-services command-list** CLI command. Dynamic Data Service SAPs must be located on dot1q- or qinq-encapsulated Ethernet, anchoring or satellite ports and can be part of a LAG.

A Python script interface adds a flexible abstraction layer that reduces the OSS integration cost; only the business user specific service parameters, such as service type, IP address, QoS, and filter parameters, are required from RADIUS or local authentication. These parameters are then used in the Python script to generate a CLI template to set up the target Dynamic Data Service.

Dynamic Data Services configuration can be updated via a RADIUS CoA message.

Both XML accounting and RADIUS accounting for up to two different RADIUS destinations can be activated on a dynamic data service SAP.

15.2 RADIUS-Triggered Dynamic Data Services Associated With a PPPoE or IPoE Session as Control Channel

Refer to the “RADIUS-Triggered Dynamic Data Service Provisioning” section in the *Advanced Configuration Guide* for a detailed description.

15.3 Data-Triggered Dynamic Data Services

In the data-triggered dynamic data services model, any frame arriving on a dynamic data services capture SAP can result in RADIUS or local authentication. The dynamic data service is then created from a Python script that generates CLI snippets using parameters obtained from authentication.

Figure 204 **Data-Triggered Dynamic Data Services**

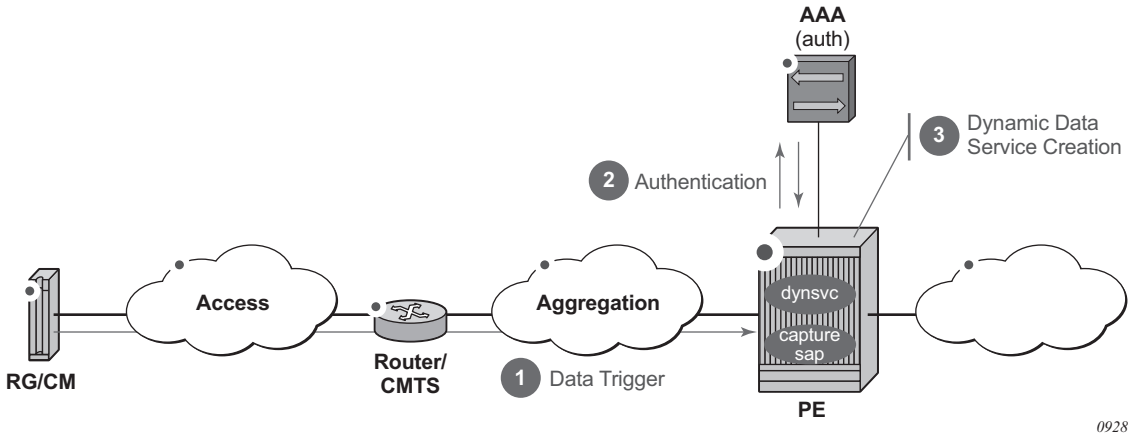

15.3.1 Data Trigger

A dynamic services data trigger is an object that is created when a frame received on a dynamic services capture SAP is sent to the control plane for authentication. A dynamic services data trigger is uniquely identified with its SAP ID. If the dynamic services data trigger was received on capture SAP x / y / z:*.* with outer-tag = a and inner-tag = b, then the dynamic service data trigger SAP ID is “x / y / z:a.b”.

For each dynamic services data trigger, the following information is stored.

Table 170 **Data Trigger Information**

Data Trigger Information	Description
Acct-Session-ID	The RADIUS accounting session ID for this dynamic services data trigger. This accounting session ID is used as accounting multi-session ID in RADIUS accounting for associated dynamic services. It can also be used as a key in CoA or Disconnect Messages to set up or terminate associated dynamic services.
MAC	The MAC address learned to set up this dynamic service data trigger. The MAC address is included in the Access-Request message for RADIUS authentication.
IP	The IPv4 or IPv6 address learned to set up this dynamic service data trigger. If the data trigger packet was not an IP packet, then this field is empty. When available, the IP address is included in the RADIUS authentication and accounting messages.
State	The current state of the dynamic service data trigger: Pending: (initial state) data trigger received and authentication started Accepted: (transient state) authentication succeeded; dynsvc script started but not yet completed sapCreated: (final state) corresponding dynamic services SAP created

The dynamic services data trigger information can be displayed as follows:

```
# show service dynamic-services data-triggers
===============================================================================
Dynamic Services Data-triggers
===============================================================================
SAP                         : 1/1/4:1214.101
-------------------------------------------------------------------------------
Acct session-ID             : 144DFF0000009156A24138
MAC                         : 00:51:00:dd:01:01
IP                          :
State                       : sapCreated
-------------------------------------------------------------------------------
No. of Data-triggers: 1
===============================================================================
```

For a data-triggered dynamic data service to be successfully set up, a dynamic services SAP equal to the data trigger SAP ID must be created.

In the same way as the control channel model, multiple dynamic data services can be associated with a single dynamic services data trigger: up to 32 dynamic services during data trigger authentication or up to 4000 in total through CoA. When the dynamic services SAP that corresponds to a data trigger is deleted (teardown action), then all dynamic services associated with that dynamic services data trigger are deleted (teardown action).

15.3.2 Dynamic Services Data Trigger Capture SAP

In the same way as an Enhanced Subscriber Management (ESM) capture SAP is configured, a dynamic services data trigger capture SAP is configured in a VPLS service and captures frames for authentication. A dynamic service data trigger capture SAP does not forward traffic within the VPLS service, and no MAC learning occurs.

A VPLS capture SAP becomes a dynamic services data trigger capture SAP when a dynamic services policy is configured and the **dynamic-services** context is enabled (**no shutdown**). The dynamic services policy specifies the authentication mechanism to be used as detailed in the authentication sections below.

For example:

```
configure
    service
        vpls 10 customer 1 create
            sap 1/1/4:1214.* capture-sap create
                description "Dynamic Services Data Trigger capture-sap"
                dynamic-services
                    dynamic-services-policy "dyn-svc-1"
                    no shutdown
                exit
                no shutdown
            exit
            no shutdown
        exit
    exit
```

The **trigger-packet** type needs to be configured. A dynamic services data trigger **capture-sap** captures any valid Ethernet frames, including non-IP frames.

Valid dynamic services data trigger capture SAPs are:

- dot1q-encapsulated Ethernet, anchoring, satellite ports or LAGs
  - sap x/y/z:*

- lag-x:*
- pxc-x.a:*
- esat-x/y/z:*

• qinq-encapsulated Ethernet, anchoring, satellite ports or LAGs:

- sap x/y/z:*.*
- sap lag-x:*.*
- sap x/y/z:tag.*
- sap x/y/z:*.tag
- sap lag-x:tag.*
- sap lag-x:*.tag
- pxc-x.a:*.*
- pxc-x.a:tag.*
- esat-x/y/z:*.*
- esat-x/y/z:tag.*

Enhanced Subscriber Management (ESM) and dynamic services data trigger cannot be enabled simultaneously on a single capture SAP: a **no shutdown** command of the **dynamic-services** CLI context is mutually exclusive when configuring an ESM trigger-packet type.

Use the following CLI commands to display capture SAP statistics:

• To display the number of data trigger packets forwarded to the CPM and dropped on the IOM:

```
# show service id <service-id> sap <sap-id> sap-stats
===============================================================================
Service Access Points(SAP)
===============================================================================
--- snip ---
-------------------------------------------------------------------------------
Sap Statistics
-------------------------------------------------------------------------------
Last Cleared Time     : N/A
                        Packets                 Octets
CPM Ingress           : 108                     6696
Forwarding Engine Stats
Dropped               : 6144                    405504
```

Table 171 **Data Trigger Packets Forwarded/Dropped Counters**

Counter	Description
CPM Ingress	Number of dynamic service data triggers received on the capture SAP that are forwarded to CPM

Table 171 **Data Trigger Packets Forwarded/Dropped Counters (Continued)**

Counter	Description
Forwarding Engine Stats Dropped	Number of dynamic service data triggers is received on the capture SAP that are dropped on IOM

- To display the dynamic services capture SAP control plane statistics (data triggers received and data trigger drop reason counters):

```
# show service id 10 dynamic-services capture-sap 1/1/4:1214.* statistics
===============================================================================
Dynamic Services Capture SAP 1/1/4:1214.* Statistics
===============================================================================
Data packets received by SAP                          : 3
-------------------------------------------------------------------------------
Drop Reason Counters
-------------------------------------------------------------------------------
No policy configured at capture SAP level             : 0
No authentication configured in policy                : 0
Data-trigger already exists                           : 0
Lockout is active                                     : 0
Reached data-trigger system limit                     : 0
No memory available                                   : 0
Unsuccessful authentication                           : 1
No data-trigger SAP-id in authentication              : 0
Corresponding dynamic SAP is not created              : 0
===============================================================================
```

Table 172 **Data Triggers Received/Dropped Reason Counters**

Counter	Description
Data packets received by SAP	Number of dynamic service data triggers received on the **capture-sap** that reached the CPM
No policy configured at capture SAP level	No dynamic services policy configured at the **capture-sap**; required to determine the authentication destination.
No authentication configured in policy	The authentication section in the specified dynamic services policy is missing or incomplete.
Data-trigger already exists	A new data trigger frame is received for an existing data trigger that is authenticated, but the corresponding dynamic SAP is not yet created. The new data trigger packet is dropped.
Lockout is active	The data trigger for this managed SAP is currently in a lockout state due to previous authentication failures.
Reached data-trigger system limit	The maximum number of dynamic service data triggers supported on the system is reached. Additional data triggers are dropped.

Table 172 **Data Triggers Received/Dropped Reason Counters (Continued)**

Counter	Description
No memory available	There is not enough system memory available to process the data trigger.
Unsuccessful authentication	The authentication for a data trigger on this capture SAP failed or timed out.
No data-trigger SAP-id in authentication	The dynamic services data trigger SAP ID is not provided in authentication. This is a mandatory parameter.
Corresponding dynamic SAP is not created	The data trigger successfully authenticated but the corresponding dynamic SAP was not created. This is typically caused by a dynamic services script error.

- to clear the dynamic services capture SAP control plane statistics:

```
# clear service id <service-id> dynamic-services capture-sap <sap-id> statistics
```

15.3.3 RADIUS Authentication

When a valid Ethernet frame is received on a dynamic services data trigger capture SAP, it is sent to the control plane for authentication. The dynamic services policy configured at the capture SAP specifies the RADIUS authentication parameters, as shown in the following example:

```
configure service
        vpls 10 customer 1 create
            sap 1/1/4:1214.* capture-sap create
                description "Dynamic Services Data Trigger capture-sap"
                dynamic-services
                    dynamic-services-policy "dyn-svc-1"
                    no shutdown
                exit
                no shutdown
            exit
            no shutdown
        exit
        dynamic-services
            dynamic-services-policy "dyn-svc-1" create
                ---snip---
                authentication
                    password "RwXx4x0jao776C3CGlDBKVaNOd//ySXw" hash2
                    server-policy "aaa-server-policy-1"
                exit
                ---snip---
            exit
        exit
```

Local authentication and RADIUS authentication are mutually exclusive and cannot be configured simultaneously in a **config>service>dynsvc>plcy>authentication** context.

The **server-policy** CLI command references the **config>aaa>radius-server-policy** *policy-name* to be used for authentication.

The **password** CLI command specifies the password that is used in all RADIUS Access-Request messages.

RADIUS Access-Request Message Attributes specifies the attributes that are included in the RADIUS Access-Request message for dynamic services data triggers.

Table 173 **RADIUS Access-Request Message Attributes**

RADIUS Attribute	Description
[1] User-Name	The username format for dynamic services data trigger authentication is fixed to *nas-port-id* (SAP).
[2] Password	The password as configured in the authentication section of the **dynamic-services-policy**.
[4] NAS-IP-Address	The outband management interface or system interface IPv4 address. Only included if the RADIUS server is reachable via an IPv4 address.
[95] NAS-IPv6-Address	The outband management interface or system interface IPv6 address. Only included if the RADIUS server is reachable via an IPv6 address.
[44] Acct-Session-Id	A unique accounting session ID (number format) per dynamic service data trigger. Included as [50] Acct-Multi-Session-Id in radius accounting for all dynamic services that are associated with this data trigger.
[87] NAS-Port-Id	The dynamic service data trigger **sap-id**
[32] NAS-Identifier	The system name of the router
[26-6527-27] Alc-Client-Hardware-Addr	The MAC address of the data trigger frame that resulted in the authentication. Fixed format (xx:xx:xx:xx:xx:xx)
[8] Framed-IP-Address	The IPv4 source address of the IPv4 data trigger frame that resulted in the authentication. Not included if the data trigger frame is not an IPv4 packet.
[26-6527-99] Alc-Ipv6-Address	The IPv6 source address of the IPv6 data trigger frame that resulted in the authentication. Not included if the data trigger frame is not an IPv6 packet.

The attributes that must be returned in the Access-Accept message are the same as for RADIUS-triggered Dynamic Data Services associated with an IPoE or PPPoE session as a control channel.

15.3.4 Local Authentication

Local authentication is available for data-triggered dynamic services deployments where RADIUS is used for accounting and dynamic changes (CoA) but cannot provide the actual service provisioning parameters.

When a valid Ethernet frame is received on a dynamic services data trigger capture SAP, it is sent to the control plane for authentication. The dynamic services policy configured at the capture SAP specifies the local authentication parameters, as shown in the following example:

```
configure service
        vpls 10 customer 1 create
            sap 1/1/4:1214.* capture-sap create
                description "Dynamic Services Data Trigger capture-sap"
                dynamic-services
                    dynamic-services-policy "dyn-svc-2"
                    no shutdown
                exit
                no shutdown
            exit
            no shutdown
        exit
        dynamic-services
            dynamic-services-policy "dyn-svc-2" create
                ---snip---
                authentication
                    local-auth-db "dynsvc-db-1"
                exit
                ---snip---
            exit
        exit
```

Local authentication and RADIUS authentication are mutually exclusive and cannot be configured simultaneously in a **config>service>dynsvc>plcy>authentication** context.

The **local-auth-db** CLI command references the local authentication database to be used for authentication, as shown in the following example:

```
configure service
        dynamic-services
            local-auth-db "dynsvc-db-1" create
                user-name "1/1/4:1214.101" create
                    description "dynsvc: epipe"
                    index 1 create
```

```
                        dynamic-services-policy "dyn-svc-2"
                        sap-id "1/1/4:1214.101"
                        script-parameters-1 "epipe_1={'t':('dynsvc-epipe-
1',None,None,10,11)}"
                        accounting 1 create
                            stats-type volume-time
                            update-interval min 30
                        exit
                    exit
                    no shutdown
                exit
                no shutdown
            exit
        exit
```

A username is used as a key for a lookup in the local authentication database. The username format for dynamic service data triggers is fixed to the SAP ID of the data trigger. For each username entry (data trigger **sap-id**), multiple dynamic service SAPs can be specified (indexes). The index enables multiple dynamic service SAPs to be associated with a single dynamic service data trigger.

The following data can be specified for each index (dynamic service SAP) in a **user-name** entry:

- dynamic service **sap-id** (mandatory)

  The dynamic service SAP ID that is created. The SAP ID of one of the indexes must match the dynamic service data trigger **sap-id**.
- dynamic-services-policy **dynsrv-policy-name** (optional)

  Specifies the policy to use for setting up the dynamic service. If not specified, the policy provisioned at the dynamic service data trigger **capture-sap** is used.
- **script-parameters** (mandatory)

  Script parameters are used as input to the dynamic data service Python script. They are specified as four strings of up to 250 characters each. The concatenation of all four script parameter strings are passed to the Python script and must be formatted as **function-key** *dictionary*. The **function-key** specifies which Python functions is called, and *dictionary* contains the actual parameters in a Python dictionary structure format. The format should match the format of the [26-6527-165] Alc-Dyn-Serv-Script-Params attribute when RADIUS authentication is used.
- accounting overrides (optional)

  For each of the two RADIUS accounting destinations, the **stats-type** and **update-interval** can be specified. These parameters override the configured value in the dynamic services policy:

- **stats-type**: specifies if dynamic service RADIUS accounting should be enabled or disabled. RADIUS accounting is enabled by specifying the statistics type: volume and time or time only. RADIUS accounting is disabled when **no stats-type** is specified.
- **update-interval**: specifies the time between each dynamic data service accounting interim update. The generation of interim accounting updates is disabled when **no update-interval** is specified.

A local authentication database can only be used to authenticate a dynamic service data trigger and provide parameters to set up associated dynamic services. The script action cannot be specified and is always set to “setup”.

The setup timeout for Access=Accept (CLI command: **configure service dynamic-services timers setup-timeout access-accept** *timeout*) also applies for local authenticated dynamic services.

15.3.5 Data-Triggered Dynamic Service Provisioning

Once authenticated, the mechanism to set up, modify, and tear down a data triggered dynamic service is the same as for RADIUS-triggered Dynamic Data Services associated with an IPoE or PPPoE session as a control channel.

The auto-provisioning of a data-triggered dynamic service is initiated by the RADIUS messages or local authentication as listed in Table 174.

Notes:

- Using the Nas-Port-Id as a key in a CoA or Disconnect-Message targets the corresponding dynamic services SAP; this also occurs when the Nas-Port-Id corresponds with the SAP ID of a data trigger.
- Using the Acct-Session-Id as a key in a CoA or Disconnect-Message targets:
  - the corresponding dynamic services SAP if the Acct-Session-Id belongs to a dynamic services SAP that is not a dynamic services data trigger SAP
  - the data trigger SAP if the Acct-Session-Id belongs to the dynamic services data trigger SAP
- In the event of a tear down, if the dynamic services SAP ID is a dynamic service data trigger SAP ID, all dynamic services associated with that dynamic services data trigger are also removed.

Table 174 **Dynamic Service Script Actions**

Action	Dynamic Service Script Action	Comments
Rx Access-Accept or local authentication (dynamic services data trigger authentication)	Setup	Up to 32 dynamic data services SAPs in a single message The dynamic services SAP that corresponds with the data trigger (also referred to as the **dynamic services data trigger sap-id**) must be part of this list. The Alc-Dyn-Serv-Script-Action VSA is optional for RADIUS authentication.
	Modify / Teardown	Not supported
Rx CoA (Nas-Port-Id or Acct-Session-Id of a dynamic service SAP different from the data trigger)	Setup	Not supported
	Modify	Only a single dynamic data service per message Mandatory VSAs: Alc-Dyn-Serv-Script-Action Alc-Dyn-Serv-Script-Params
	Teardown	Tear down the dynamic service of the dynamic services SAP identified by the Acct-Session-Id or Nas-Port-Id. Alc-Dyn-Serv-Script-Action VSA is mandatory
Rx CoA (Nas-Port-Id of a data trigger)	Setup	Not supported. Nas-Port-Id always targets the dynamic services SAP and not the data trigger.
	Modify	Only a single dynamic data service per message Mandatory VSAs: Alc-Dyn-Serv-Script-Action Alc-Dyn-Serv-Script-Params
	Teardown	Tear down the dynamic service of the dynamic services SAP identified by the Nas-Port-Id. Because this is the data trigger SAP that is deleted, all dynamic services SAPs associated with the data trigger are also deleted. Alc-Dyn-Serv-Script-Action VSA is mandatory

Table 174 **Dynamic Service Script Actions (Continued)**

Action	Dynamic Service Script Action	Comments
Rx CoA (Acct-Session-Id of a data trigger)	Setup	Only a single dynamic service SAP per message When successful, the dynamic services SAP is associated with the data trigger identified by the specified Acct-Session-Id. Mandatory VSAs: Alc-Dyn-Serv-Script-Action Alc-Dyn-Serv-SAP-Id Alc-Dyn-Serv-Script-Params Alc-Dyn-Serv-Policy (if no “default” policy configured)
	Modify	Only a single dynamic service per message Modify the dynamic service of the dynamic services SAP identified by the Alc-Dyn-Serv-SAP-Id. The dynamic services SAP must be associated with the data trigger identified with the specified Acct-Session-Id. Mandatory VSAs: Alc-Dyn-Serv-Script-Action Alc-Dyn-Serv-SAP-Id Alc-Dyn-Serv-Script-Params
	Teardown	Tear down the dynamic service of the dynamic services SAP identified by the Alc-Dyn-Serv-SAP-Id. The dynamic services SAP must be associated with the data trigger identified with the specified Acct-Session-Id. If the dynamic services SAP identified by the Alc-Dyn-Serv-SAP-Id is a data trigger sap, then teardown the dynamic services of all dynamic services saps associated with that data trigger Mandatory VSAs: Alc-Dyn-Serv-Script-Action Alc-Dyn-Serv-SAP-Id
Rx Disconnect Message (Nas-Port-Id or Acct-Session-Id of a dynamic service SAP different from a data trigger)	N/A	Tear down the dynamic service of the dynamic services SAP identified by the Acct-Session-Id or Nas-Port-Id

Table 174 **Dynamic Service Script Actions (Continued)**

Action	Dynamic Service Script Action	Comments
Rx Disconnect Message (Nas-Port-Id or Acct-Session-Id of a data trigger)	N/A	Tear down the dynamic services of all dynamic services SAPs associated with the data trigger identified by the Acct-Session-Id or Nas-Port-Id

A data-triggered dynamic service must be explicitly removed by one of the following:

- with a RADIUS Disconnect message containing the Acct-Session-Id or NAS-Port-Id as key
- with a RADIUS CoA message containing the Acct-Session-Id or NAS-Port-Id as key and Alc-Dyn-Serv-Script-Action VSA with value 3 (teardown)
- with a CLI **clear** command: **clear>service>dynamic-services>data-trigger sap** *sap-id*

  All dynamic service SAPs associated with the dynamic services data trigger is removed.
- with a CLI tools command: **tools>perform>service>dynamic-services> evaluate-script sap** *sap-id* **control-session** *acct-session-id* **action teardown**

  The control session accounting session ID corresponds to the dynamic services data trigger accounting session ID.

The removal of a dynamic service SAP that is a data trigger SAP results in the removal (teardown) of all dynamic service SAPs associated with that dynamic services data trigger.

To prevent a data-triggered dynamic service from being immediately set up again after it was removed (because traffic is still being received), the following procedures can be used:

- Authentication failure
  - Update the configuration of the RADIUS server or local authentication such that the authentication for the dynamic service data trigger fails
  - Tear down the dynamic service
  - The dynamic service is not set up again because the data trigger authentication fails, resulting in a host-lockout when provisioned.
- VID filter on the **capture-sap**

- Add the data trigger encapsulation to a VID filter (ingress MAC filter of type **vid**) that is applied on the data trigger **capture-sap**
- Tear down the dynamic service
- The dynamic service is not set up again because the data trigger is now dropped by the VID filter applied on the **capture-sap**

15.3.6 Control Plane Protection

As a dynamic services data trigger **capture-saps** potentially forward all valid Ethernet frames for authentication to the control plane, control plane protection mechanisms are required to prevent overload conditions.

1. **capture-sap** data trigger packet throttling (frames dropped at IOM)

   The number of data trigger packets sent to the control plane via the ingress forwarding complex is rate-limited based on a hash using the sap ID, outer tag, and inner tag as the key. The per-hash result is that a maximum of 1 frame is forwarded to the control plane.

   This throttling mechanism is always enabled and has no configuration options. It guarantees fairness between different encapsulation while limiting the frame rate sent to the control plane.

2. Blocking certain VLANs from authenticating (frames dropped at the IOM)

   This can be achieved by applying ingress MAC filters of type VID with the **capture-sap** command. In the example below, frames with encap 1/1/4:1214.20 is dropped by the VID filter.

```
configure
    service
        vpls <service-id> customer <customer-id>
            sap 1/1/4:1214.* capture-sap
                dynamic-services
                    dynamic-services-policy <dynsvc-policy-name>
                    no shutdown
                exit
                ingress
                    filter mac 10
                exit
            exit
        exit
    exit
    filter
        mac-filter 10 create
            default-action forward
            type vid
            entry 10 create
                match frame-type ethernet_II
                    outer-tag 20 4095
                exit
```

```
            action
                drop
            exit
        exit
    exit
exit
```

3. Dynamic service data trigger rate limiting in the control plane (frames dropped at the CPM)

   An overall rate limit of dynamic service data triggers limits the number of frames that come from the different IOMs to an acceptable rate for the control plane to handle.

   Control plane rate-limiting for dynamic service data triggers is always enabled and has no configuration options.

4. Using a lockout mechanism to protect the RADIUS infrastructure from overload as a result of permanent failing authentications of Python script errors.

   The existing Enhanced Subscriber Management (ESM) host lockout mechanism can be enabled on a dynamic services data trigger **capture-sap**.

   The dynamic services data trigger **sap-id** is used as a key for the **host-lockout** context. A **host-lockout** context is created whenever a data trigger is deleted:

   – Authentication failures: RADIUS Access Reject, RADIUS Access-Accept with a wrong or missing data trigger SAP ID, timeout, local authentication lookup failure, local authentication returns a wrong or **no data-trigger sap-id**.

   – No data trigger SAP created within the configured dynamic services **access-accept setup timeout** (30 seconds by default):

     ```
     configure service dynamic-services timers setup-timeout access-accept <timeout>
     ```

   – A dynamic services Python script failure

   – A **clear** command on a data trigger:

     ```
     clear service dynamic-services data-trigger sap <sap-id>
     ```

   – A tear down of a data trigger initiated via a CoA or Disconnect Message

   The last two bullets are tear down operations that occur infrequently and should not result in an actual lockout; only the host lockout context is created and should disappear again when the **lockout-reset-time** expires. This is configured in the **host-lockout-policy**:

```
configure subscriber-mgmt host-lockout-policy <policy-name> lockout-reset-time
<seconds>   (default 60 seconds).
```

   To enable host lockout for dynamic services data trigger, configure a **host-lockout-policy** at the dynamic services data trigger **capture-sap**:

```
configure service
```

```
vpls <service-id> customer <customer-id>
    sap <sap-id> capture-sap
        dynamic-services
            dynamic-services-policy <dynsvc-policy-name>
            no shutdown
        exit
        host-lockout-policy <policy-name>
    exit
exit
```

The **host-lockout-policy** is configured in the **config>subscr-mgmt** CLI context. For data-triggered dynamic services, the **host-key** in the policy must be set to **all** which is achieved with the **no host-key** CLI command (default). Configuring the host key to **mac** in a host lockout policy associated with a dynamic service data trigger **capture-sap** is a configuration error and results in a faulty host lockout behavior.

Use the following show command to display active dynamic services data trigger capture SAP lockouts (the capture SAP must be used as the **sap-id**):

```
# show subscriber-mgmt host-lockout-policy "host-lockout-1" all sap 1/1/4:1214.*
===============================================================================
Host Lockout Policy "host-lockout-1"
===============================================================================
Description                     Host lockout policy
Last Mgmt Change                01/25/2016 16:37:47
Lockout time min                10
Lockout time max                3600
Lockout reset time              60
Max lockout hosts               100
Host key                        all
-------------------------------------------------------------------------------
Active Lockouts for SAP: 1/1/4:1214.*
===============================================================================
circuit-id/                     elapsed  current  elapsed  next     nr
  mac/                          reset    lock     lock     lock     of
remote-id                       time (s) time (s) time (s) time (s) lockouts
-------------------------------------------------------------------------------
:1214.101                       0        10       1        20       1
-------------------------------------------------------------------------------
Nr of active lockouts           1
Nr of lockouts in grace period  0
Nr of total lockouts:           1
-------------------------------------------------------------------------------
Totals for Host Lockout Policy "host-lockout-1"
-------------------------------------------------------------------------------
Nr of active lockouts           1
Nr of lockouts in grace period  0
Nr of total lockouts:           1
===============================================================================
```

Use following command to clear active dynamic services data trigger capture SAP lockouts (the capture SAP must be used as the **sap-id**):

```
# clear subscriber-mgmt host-lockout-policy [policy <host-lockout-policy-
name>] <lockout-state>
# clear subscriber-mgmt host-lockout-policy sap <sap-id> [lockout-state]
```

5. Use the **cpu-protection** and **dist-cpu-protection** commands on the platforms that support it.

   When using the **cpu-protection** command, a **cpu-protection** *policy-id* is needed to apply on the capture SAP.

```
configure
    system
        security
            cpu-protection
                policy 200 create
                    overall-rate 100
                    out-profile-rate 50
                exit
            exit
        exit
    exit
    service
        vpls <service-id> customer <customer-id>
            sap 1/1/1:*.* capture-sap create
                cpu-protection 200
                dynamic-services
                    dynamic-services-policy "dynServPolicy"
                    no shutdown
                exit
                no shutdown
            exit
        exit
    exit
```

   When using the **dist-cpu-protection** command, a **dist-cpu-protection** *policy-id* is needed to apply on the capture SAP.

```
configure
    system
        security
            dist-cpu-protection
                policy "distCpuProt" create
                    static-policer "dcpuStatPol" create
                        rate packets 50 within 1
                        exceed-action discard
                        detection-time 5
                    exit
                    protocol all-unspecified create
                        enforcement static "dcpuStatPol"
                    exit
                exit
            exit
        exit
    exit
    service
        vpls <service-id> customer <customer-id>
            sap 1/1/1:*.* capture-sap create
                dynamic-services
                    dynamic-services-policy "dynServPolicy"
                    no shutdown
                exit
                dist-cpu-protection "distCpuProt"
                no shutdown
```

```
            exit
        exit
    exit
```

15.3.7 Debugging

To enable dynamic services data trigger debugging, use the following debug commands:

```
debug
    dynamic-services
        data-triggers
            # multiple capture-sap commands allowed
            capture-sap <sap-id> [encap-val <qtag[.qtag]>] [mode <mode>]
            capture-sap <sap-id> [encap-val <qtag[.qtag]>] [mode <mode>]
        exit
    exit
```

During debugging, the system logs data trigger events such as:

- data trigger received
- data trigger authentication events
- data trigger SAP created
- dynamic service SAP created
- dropped data trigger with drop reason such as data trigger already exists or lockout active

The **encap-val** command limits the debug output to data trigger events for specific encapsulation values.

The mode specifies which data trigger events are logged: all events or dropped data trigger events only.

15.4 Dynamic Data Services Python API

Table 175 describes the functions available in the SR OS alc.dyn Python module to create dynamic data services.

Table 175 **Dynamic Data Services Functions in the alc.dyn Module**

alc.dyn functions	Description
dyn.action (dictionary)	Executes the setup, modify or teardown function that is found with a lookup in the specified dictionary using the function-key present in the script parameters. The script-parameters can be obtained from RADIUS in the Alc-Dyn-Serv-Script-Params VSA or can be configured in the local authentication database. The action (setup, modify or teardown) is determined from the Alc-Dyn-Script-Action VSA. Returns: None (no value) *dictionary* - Python dictionary with format: dictionary = {function-key-1 : (setup-function-1, modify-function-1|None, revert-function-1|None, teardown-function-1), …, function-key-n : (setup-function-n, modify-function-n|None, revert-function-n|None, teardown-function-n) } where: • setup-function-n: name of the function invoked when Alc-Dyn-Serv-Script-Action=setup • modify-function-n: name of the function invoked when Alc-Dyn-Serv-Script-Action=modify. The modify function is optional. When not specified, the keyword “None” should be used in the dictionary. • revert-function-n: name of the function invoked when modify script execution failed. This function must undo all changes from the modify function. The revert function is optional. When not specified, the keyword “None” should be used in the dictionary. • teardown-function-n: name of the function invoked when Alc-Dyn-Serv-Script-Action=teardown
dyn.add_cli (string)	Executes the specified CLI commands to create the dynamic data service Returns: None (no value) *string* - CLI commands. The string can span multiple lines in the script when enclosed in three double quotes ("""").
dyn.get_circuit-id()	Returns a string containing the control channel circuit-id (DHCP relay agent option 82 or PPP tags)

Table 175 **Dynamic Data Services Functions in the alc.dyn Module (Continued)**

alc.dyn functions	Description
dyn.get_remote-id()	Returns a string containing the control channel remote-id (DHCP relay agent option 82 or PPP tags)
dyn.get_sap()	Returns a string containing the dynamic data service sap id: the value of the Alc-DynServ-SAP-Id VSA. A wildcard hash (#) in the VSA value is replaced with the corresponding control channel or data trigger port or vlan id field.
dyn.reference (function-key, reference-id, dictionary)	Create a dynamic reference to another function in the script. Typically used to create n:1 relations between dynamic data services, such as multiple SAPs in a VPLS service or multiple services for the same customer. Dynamic referencing is only allowed in setup functions. Corresponding teardown functions will automatically dereference. Returns a dictionary provided by the setup script matching the function-key: *function-key* - key in the action dictionary (see dyn.action) to find the corresponding setup or teardown functions. *reference-id* - unique reference ID string that identifies the dynamic reference (for example, all SAPs from a VPLS service would have the same reference id). *dictionary* - a Python dictionary containing parameters for use in the setup or teardown function, such as to generate CLI output.
dyn.select_free_id (type)	Returns a string representing a free identifier of the specified type. *type* - identifier type that will be returned: • "*service-id*" - obtain a free service id in the configured dynamic-services service-range • "*ip-filter-id*" - generate a new ip-filter identifier • "*ipv6-filter-id*" - generate a new ipv6-filter identifier • "*mac-filter-id*" - generate a new mac-filter identifier The corresponding filters will be shown in the system as "_tmnx_dyn_<number>"

15.5 RADIUS Triggered Dynamic Data Services Command Reference

15.5.1 Command Hierarchies

15.5.1.1 RADIUS Triggered Dynamic Data Services Commands

```
config
    — service
        — dynamic-services
            — dynamic-services-policy dynsrv-policy-name [create]
            — no dynamic-services-policy dynsrv-policy-name
                — accounting-1
                    — server-policy policy-name
                    — no server-policy
                    — stats-type {time | volume-time}
                    — no stats-type
                    — update-interval [hrs hours] [min minutes] [days days]
                    — no update-interval
                    — update-interval-jitter absolute seconds
                    — no update-interval-jitter
                — accounting-2
                    — server-policy policy-name
                    — no server-policy
                    — stats-type {time | volume-time}
                    — no stats-type
                    — update-interval [hrs hours] [min minutes] [days days]
                    — no update-interval
                    — update-interval-jitter absolute seconds
                    — no update-interval-jitter
                — authentication
                    — local-auth-db name
                    — no local-auth-db
                    — password password [hash | hash2]
                    — no password
                    — server-policy policy-name
                    — no server-policy
                — cli-user name
                — no cli-user
                — description description-string
                — no description
                — sap-limit [limit]
                — no sap-limit
                — script-policy name
                — no script-policy
            — local-auth-db name [create]
```

```
— no local-auth-db name
    — description description-string
    — no description
    — [no] shutdown
    — user-name name [create]
    — no user-name name
        — description description-string
        — no description
        — index index [create]
        — no index index
            — accounting {1 | 2} [create]
            — no accounting {1 | 2}
                — stats-type {volume-time | time}
                — no stats-type
                — update-interval [hrs hours] [min minutes] [days
                    days]
                — no update-interval
            — dynamic-services-policy name
            — no dynamic-services-policy
            — sap-id sap-string
            — no sap-id
            — script-parameters-1 param-string1
            — no script-parameters-1
            — script-parameters-2 param-string2
            — no script-parameters-2
            — script-parameters-3 param-string3
            — no script-parameters-3
            — script-parameters-4 param-string4
            — no script-parameters-4
        — [no] shutdown
— service-range service-id service-id
— no service-range
— timers
    — setup-timeout access-accept timeout
    — no setup-timeout
```

```
configure
    — service
        — vpls service-id customer customer-id [create]
            — sap sap-id capture-sap [create]
                — dynamic-services
                    — dynamic-services-policy name
                    — no dynamic-services-policy
                    — [no] shutdown
                — host-lockout-policy policy-name
                — no host-lockout-policy
```

```
config
    — system
```

```
            — security
                — password
                    — dynsvc-password password [hash | hash2]
                    — no dynsvc-password

<global>
    — [no] enable-dynamic-services-config
```

15.5.1.2 Basic System Command

```
<global>
    — info
```

15.5.1.3 Show Commands

```
show
    — service
        — dynamic-services
            — data-triggers [sap sap-id] [summary]
            — dynamic-services-policy [policy-name]
            — root-objects
            — saps [control-session acct-session-id] [port port-id] [dynsvc-policy policy-
                name] [summary] [orphaned] [sap sap-id] [svc-id service-id]
            — script
                — snippets [detail]
                — snippets name snippet-name [instance snippet-instance] [detail]
                — statistics
            — summary
        — id service-id
            — dynamic-services
                — capture-sap [statistics] sap-id
        — sap-using [msap] [dyn-script] [description]
```

15.5.1.4 Clear Commands

```
clear
    — service
        — dynamic-services
            — data-trigger sap sap-id
        — statistics
            — dynamic-services
```

15.5.1.5 Debug Commands

```
debug
    — [no] dynamic-services
        — data-triggers
        — no data-triggers
            — capture-sap sap-id [encap-val qtag[.qtag]] [mode mode]
            — no capture-sap sap-id
        — scripts
            — [no] event
                — [no] cli
                — [no] errors
                — [no] executed-cmd
                — [no] state-change
                — [no] warnings
            — instance instance
                — [no] event
                    — [no] cli
                    — [no] errors
                    — [no] executed-cmd
                    — [no] state-change
                    — [no] warnings
            — script script
                — [no] event
                    — [no] cli
                    — [no] errors
                    — [no] executed-cmd
                    — [no] state-change
                    — [no] warnings
```

15.5.1.6 Tools Commands

```
tools
    — perform
        — service
            — dynamic-services
                — evaluate-script sap sap-id control-session acct-session-id action
                  script-action [dynsvc-policy name] [param-string string]
    — dump
        — service
            — dynamic-services command-list
```

15.5.2 RADIUS Triggered Dynamic Data Services Commands

15.5.2.1 Configuration Commands

dynamic-services

Syntax **dynamic-services**

Context config>service>vpls>sap

Description This command enables the context to configure dynamic services parameters on a capture SAP.

shutdown

Syntax [**no**] **shutdown**

Context config>service>vpls>sap>dyn-svc

Description This command disables or enables data-triggered dynamic services on this **capture-sap**.

dynamic-services

Syntax **dynamic-services**

Context config>service

Description This command enables the context to configure dynamic data services. Only available on systems with multi-core CPM (CPM3 or up).

dynamic-services-policy

Syntax **dynamic-services-policy** *dynsrv-policy-name* [**create**]
no dynamic-services-policy *dynsrv-policy-name*

Context config>service>dynsvc

Description This command creates a new dynamic services policy that can be used to create dynamic data services.

The **no** form of this command removes the dynamic services policy from the configuration. This is only allowed when there are no active dynamic data services referencing this policy.

Parameters *dynsrv-policy-name* — Specifies a unique name of a dynamic services policy up to 32 characters.

accounting-1

Syntax **accounting-1**

Context config>service>dynsvc>policy

Description This command enables the context to configure the first RADIUS accounting destination and corresponding RADIUS accounting parameters for dynamic data services.

accounting-2

Syntax **accounting-2**

Context config>service>dynsvc>policy

Description This command enables the context to configure the second RADIUS accounting destination and corresponding RADIUS accounting parameters for dynamic data services.

server-policy

Syntax **server-policy** *policy-name*
no server-policy

Context config>service>dynsvc>acct-1
config>service>dynsvc>acct-2

Description This command configures the radius server policy to be used for dynamic data services RADIUS accounting.

The **no** form of this command removes the radius server policy from the configuration. This is only allowed when there are no active dynamic data services referencing this policy.

Parameters *policy-name* — Specifies the name of the radius server policy.

Values Up to 32 characters maximum

stats-type

Syntax **stats-type** {**time** | **volume-time**}

no stats-type

Context config>service>dynsvc>acct-1
config>service>dynsvc>acct-2

Description This command configures the type of statistics to be reported in dynamic data services RADIUS accounting. A RADIUS specified Stats Type overrides the CLI configured value.

The **no** form of this command resets the default value.

Default stats-type volume-time

Parameters *time* — Only report Session-Time in the RADIUS Accounting Interim-Update and Stop message.

volume-time — Report both Session-Time and Volume counter attributes in the RADIUS. Accounting Interim-Update and Stop messages.

update-interval

Syntax **update-interval** [**hrs** *hours*] [**min** *minutes*] [**days** *days*]
no update-interval

Context config>service>dynsvc>acct-1
config>service>dynsvc>acct-2

Description This command specifies the interval between each RADIUS Accounting Interim-Update message (minimum 5 minutes; maximum 180 days).

The **no** form of this command disables the sending of Accounting Interim-Update messages.

A RADIUS specified Accounting Interim Interval overrides the CLI configured value.

Default no update-interval (do not send Accounting Interim-Update messages)

Parameters *hrs* — specifies the interval in hours.

Values 1 to 23

min — Specifies the interval in minutes.

Values 1 to 59

days — specifies the interval in days.

Values 1 to 180

update-interval-jitter

Syntax **update-interval-jitter absolute** *seconds*
no update-interval-jitter

Context config>service>dynsvc>acct-1
config>service>dynsvc>acct-2

Description This command specifies the absolute maximum random delay introduced on the update interval between two RADIUS Accounting Interim Update messages. The effective maximum random delay value is the minimum of the configured absolute jitter value and 10% of the configured update-interval.

A value of zero sends the accounting interim update message without introducing an additional random delay.

The **no** form of this command sets the default to 10% of the configured update-interval.

Default no update-interval-jitter (10% of the configured update-interval)

Parameters *seconds* — Specifies the absolute maximum jitter value in seconds.

Values 0 to 3600

authentication

Syntax **authentication**

Context config>service>dynsvc>policy

Description This command enables the context to configure authentication parameters for data-triggered dynamic services.

local-auth-db

Syntax **local-auth-db** *name*
no local-auth-db

Context config>service>dynsvc>policy>auth

Description This command configures the local authentication database to be used for local authentication of data-triggered dynamic services.

Local authentication and RADIUS authentication are mutually exclusive.

The **no** form of this command removes the local authentication database from the configuration and disables local authentication.

Parameters *name* — local authentication database name, up to 32 characters

password

Syntax **password** *password* [**hash** | **hash2**]
no password

Context config>service>dynsvc>policy>authentication

Description This command configures the password to be used for RADIUS authentication of data-triggered dynamic services.

The **no** form of this command removes the password from the configuration.

Parameters *password* — Specifies the password that is used in RADIUS authentication of a data-triggered dynamic service. The maximum length is 20 characters if unhashed, 32 characters if hashed, and 54 characters if the **hash2** keyword is specified.

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

server-policy

Syntax **server-policy** *policy-name*
no server-policy

Context config>service>dynsvc>policy>authentication

Description This command configures the RADIUS server policy to be used for RADIUS authentication of data-triggered dynamic services.

Local authentication and RADIUS authentication are mutually exclusive.

The **no** form of this command removes the server policy from the configuration and disables RADIUS authentication.

Parameters *policy-name* — Specifies a RADIUS server policy name, up to 32 characters.

local-auth-db

Syntax **local-auth-db** *name* [**create**]
no local-auth-db *name*

Context config>service>dynsvc

Description This command creates a local authentication database that can be used for local authentication of data-triggered dynamic services.

The **no** form of this command removes the local authentication database from the configuration.

Parameters *name* — Specifies a local authentication database name, up to 32 characters.

shutdown

Syntax [**no**] **shutdown**

Context config>service>dynsvc>ladb

Description This command disables or enables the local authentication database. When disabled, the database cannot be used for authentication.

user-name

Syntax **user-name** *name* [**create**]
no user-name *name*

Context config>service>dynsvc>ladb

Description This command creates a user name entry in the local authentication database. The user name entry is used to match with the user name of a local authenticated dynamic service data trigger. The user name of a dynamic service data trigger is fixed to the **sap-id**. When matched, the corresponding authentication data is used to set up the dynamic data services.

The **no** form of this command removes the user name entry from the local authentication database configuration.

Parameters *name* — Specifies the user name entry name, up to 64 characters.

shutdown

Syntax [**no**] **shutdown**

Context config>service>dynsvc>ladb>user

Description This command disables or enables a user name entry in the local authentication database. When disabled, the entry is not matched.

index

Syntax **index** *index* [**create**]
no index *index*

Context config>service>dynsvc>ladb>user

Description This command creates an index entry containing authentication data for a dynamic service SAP. Up to 32 indexes can be created per user name entry, representing up to 32 dynamic service SAPs that can be instantiated with a single dynamic service data trigger. One of the dynamic service SAPs must be the data trigger SAP.

The **no** form of this command removes the index entry from the user name entry in the local authentication database configuration.

Parameters *index* — Specifies the index entry identifier.

Values 1 to 32

accounting

Syntax **accounting** {**1** | **2**} [**create**]
no accounting {**1** | **2**}

Context config>service>dynsvc>ladb>user>idx

Description This command creates a context for one of the two accounting destinations specified in the dynamic services policy. In this context, overrides of RADIUS accounting parameters can be specified.

The **no** form of this command removes the RADIUS accounting overrides context from the configuration

Parameters **{1 | 2}** — Indicates one of the two RADIUS accounting destinations

stats-type

Syntax **stats-type** {**volume-time** | **time**}
no stats-type

Context config>service>dynsvc>ladb>user>idx>acct

Description This command specifies whether dynamic service accounting should be enabled or disabled for this destination. RADIUS accounting is enabled by specifying the stats type: volume and time or time only. This command overrides the local configured value in the dynamic services policy.

The **no** form of this command disables RADIUS accounting (**stats-type** off).

Parameters **volume-time | time** — Enables RADIUS accounting for this dynamic service and specifies if volume counters should be included (**volume-time**) or time only (**time**) in the RADIUS accounting messages.

update-interval

Syntax **update-interval** [**hrs** *hours*] [**min** *minutes*] [**days** *days*]
no update-interval

Context config>service>dynsvc>ladb>user>idx>acct

Description This command specifies the time between each dynamic data service accounting interim update for this accounting destination. This command overrides the local configured value in the dynamic services policy.

The **no** form of this command disables the generation of interim accounting updates to this destination.

The minimum update interval is 5 minutes.

Parameters *hours* — Specifies the interval in hours.

Values 1 to 23

minutes — Specifies the interval in minutes.

Values 1 to 59

days — Specifies the interval in days.

Values 1 to 180

dynamic-services-policy

Syntax **dynamic-services-policy** *name*
no dynamic-services-policy

Context config>service>dynsvc>ladb>user>idx
config>service>vpls>sap>dyn-svc

Description This command specifies the local configured dynamic data service policy to use for provisioning (local authentication database context) or authentication (**capture-sap** context) of this dynamic service. If not specified, the dynamic services policy with the name **default** is used. If the default policy does not exist, then the dynamic data service setup or authentication fails.

The **no** form of this command removes the dynamic services policy from the configuration.

Parameters *name* — Specifies a dynamic services policy name, up to 32 characters.

host-lockout-policy

Syntax **host-lockout-policy** *policy-name*
no host-lockout-policy

Context config>service>vpls>sap

Description This command selects an existing host lockout policy. The **host-lockout-policy** *policy-name* is created in the **config>subscr-mgmt** context.

The **no** form of the command removes the policy name from the SAP configuration.

Parameters *policy-name* — Specifies an existing host lockout policy, up to 32 characters, to associate with the SAP.

sap-id

Syntax **sap-id** *sap-string*
no sap-id

Context config>service>dynsvc>ladb>user>idx

Description This command specifies the dynamic data service SAP that is created. A dynamic service SAP ID uniquely identifies a dynamic data service instance. For a local authenticated dynamic service data trigger, one of the dynamic service SAP IDs must be the data trigger SAP.

The **no** form of this command removes the **sap-id** from the configuration.

Parameters *sap-string* — Specifies a string representing the dynamic service SAP ID (only SAPs on Ethernet ports and LAGs are valid), up to 64 characters

script-parameters-1

Syntax **script-parameters-1** *param-string1*
no script-parameters-1

Context config>service>dynsvc>ladb>user>idx

Description This command specifies the first part of parameters as input to the dynamic data service Python script. The concatenation of all four script-parameters strings are passed to the Python script and must be formatted as function-key <dictionary>. The function-key specifies which Python functions is called, and <dictionary> contains the actual parameters in a Python dictionary structure format.

The **no** form of this command removes **script-parameters-1** from the configuration.

Parameters *param-string1* — Specifies a string representing parameters that are used as input for the dynamic service Python script, up to 250 characters

script-parameters-2

Syntax **script-parameters-2** *param-string2*
no script-parameters-2

Context config>service>dynsvc>ladb>user>idx

Description This command specifies the second part of parameters as input to the dynamic data service Python script. The concatenation of all four script-parameters strings are passed to the Python script and must be formatted as function-key <dictionary>. The function-key specifies which Python functions is called, and <dictionary> contains the actual parameters in a Python dictionary structure format.

The **no** form of this command removes the **script-parameters-2** from the configuration.

Parameters *param-string2* — Specifies a string representing parameters that are used as input for the dynamic service Python script, up to 250 characters.

script-parameters-3

Syntax **script-parameters-3** *param-string3*
no script-parameters-3

Context config>service>dynsvc>ladb>user>idx

Description This command specifies the third part of parameters as input to the dynamic data service Python script. The concatenation of all four script-parameters strings are passed to the Python script and must be formatted as function-key <dictionary>. The function-key specifies which Python functions is called, and <dictionary> contains the actual parameters in a Python dictionary structure format.

The **no** form of this command removes the **script-parameters-3** from the configuration

Parameters *param-string3* — Specifies string representing parameters that are used as input for the dynamic service Python script, up to 250 characters.

script-parameters-4

Syntax **script-parameters-4** *param-string4*
no script-parameters-4

Context config>service>dynsvc>ladb>user>idx

Description This command specifies the fourth part of parameters as input to the dynamic data service Python script. The concatenation of all four script-parameters strings are passed to the Python script and must be formatted as function-key <dictionary>. The function-key specifies which Python functions is called, and <dictionary> contains the actual parameters in a Python dictionary structure format.

The **no** form of this command removes the **script-parameters-4** from the configuration.

Parameters *param-string4* — Specifies a string representing parameters that are used as input for the dynamic service Python script, up to 250 characters.

cli-user

Syntax **cli-user** *name*
no cli-user

Context config>service>dynsvc>policy

Description This command specifies the CLI user to be used to execute the dynamic data services CLI scripts. With the specified user's profile, it is possible to further restrict the internal list of allowed commands to be executed via dynamic data service CLI scripts.

The **no** form of this command sets the CLI user to an internal user with all configuration rights.

Parameters *name* — Specifies the CLI user name that must exist in the >config>system>security CLI context.

description

Syntax **description** *description-string*
no description

Context config>service>dynsvc>policy
config>service>dynsvc>ladb
config>service>dynsvc>ladb>user

Description This command associates a text string with a configuration context to help identify the content in the configuration file.

The **no** form of this command removes the string from the configuration.

Parameters *description-string* — The description character string. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

sap-limit

Syntax **sap-limit** [*limit*]
no sap-limit

Context config>service>dynsvc>policy

Description This command specifies a limit for the number of dynamic data service instances (SAPs) that can be setup simultaneously using a given dynamic services policy.

A value of zero (0) means the policy is drained: existing dynamic data services can be modified and torn down but no new dynamic data services can be setup.

Default sap-limit 1

Parameters [0 to 131072] — Specifies the number of dynamic data service SAPs that can be setup simultaneously using this dynamic services policy.

Values 0 to 131072

script-policy

Syntax **script-policy** *name*
no script-policy

Context config>service>dynsvc>policy

Description This command specifies the radius script policy to be used to setup the dynamic data services. The script-policy configuration cannot be changed when there are active dynamic data services referencing the policy.

The **no** form of this command removes the script-policy from the configuration. This is only allowed when there are no active dynamic data services referencing this policy.

Parameters *name* — Specifies the RADIUS script policy name.

service-range

Syntax **service-range** *service-id service-id*
no service-range

Context config>service>dynsvc

Description This command specifies the service ID range that is reserved for dynamic data service creation. The range cannot overlap with existing static configured services. Once configured with active dynamic services in the range, the service range can only be extended at the end.

The **no** form of this command removes the service-range from the configuration. This is only allowed when there are no active dynamic data services.

When **no service-range** is specified, the setup of dynamic data services fails.

Parameters *service-id* — Specifies the start and end service IDs to define the service range for dynamic services.

Values 1 to 2147483647

timers

Syntax **timers**

Context config>service>dynsvc

Description This command enables the context to configure dynamic data services related timers.

setup-timeout

Syntax **setup-timeout access-accept** *timeout*
no setup-timeout

Context config>service>dynsvc>timers

Description This command specifies the time that dynamic data services setup requests from a RADIUS Access-Accept are hold in an internal work queue waiting to be processed. If after the timeout, the dynamic data service setup request is still in the queue (meaning it is not setup), then the dynamic service setup request is removed from the queue and the setup fails.

The **no** form of this command reverts to the default.

Default no setup-timeout (30 seconds)

Parameters *timeout* — Specifies the setup-timeout, in seconds, for setup requests of dynamic services received via Access-Accept.

Values 2 to 3600

dynsvc-password

Syntax **dynsvc-password** *password* [**hash** | **hash2**]
no dynsvc-password

Context config>system>security>password

Description

Note: See also the description for the **enable-dynamic-services-config** command.

This command allows a user with admin permissions to configure a system wide password which enables a user to enter a special dynamic services configuration mode.

The minimum length of the password is determined by the **minimum-length** command. The complexity requirements for the password are determined by the **complexity** command.

The **no** form of this command removes the dynsvc password from the configuration

Parameters *password* — Configures the password which enables a user to enter a special dynamic services configuration mode. The maximum length can be up to 20 characters if unhashed, 32 characters if hashed, 54 characters if the **hash2** keyword is specified.

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

enable-dynamic-services-config

Syntax [**no**] **enable-dynamic-services-config**

Context <global>

Description

Note: See also the description for the **dynsvc-password** command.

If the dynsvc-password is configured in the **config>system>security>password** context, then any user can enter a special dynamic services configuration mode by entering the **enable-dynamic-services-config** command.

The **enable-dynamic-services-config** command is not in the default profile. To give access to this command, the user must belong to the administrative profile or a new profile should be created.

Once the **enable-dynamic-services-config** command is entered, the user is prompted for a password. If the password matches, the user is given access to the dynamic services configuration. Access to static configuration is in this case prohibited.

To verify that a user is in the **enable-dynamic-services-config** mode, use the **show users** command. Users in the **enable-dynamic-services-config** mode lists the letter “D” next to the user’s CLI session.

The **no** form of this command disables the dynamic services configuration mode for this user.

15.5.2.2 Basic System Command

info

Syntax **info** [**detail**] [**objective**]

Context any configuration context

Description This command displays the running configuration for the configuration context where it is entered and all branches below that context level.

By default, the command only enters the configuration parameters that vary from the default values.

The **detail** keyword causes all configuration parameters to be displayed. The **include-dynamic** objective keyword includes configuration parameters from dynamic sources such as VSD or dynamic data services Python scripts. These dynamic configuration parameters are not saved in the configuration file.

Parameters **detail** — Displays all configuration parameters including parameters at their default values.

objective — Provides an output objective that controls the configuration parameters to be displayed.

Values **include-dynamic**: includes configuration parameters from dynamic sources such as VSD or dynamic data services Python scripts.

15.5.2.3 Show Commands

dynamic-services

Syntax **dynamic-services**

Context show>service
show>service>id

Description This command enables the context to show dynamic services information.

data-triggers

Syntax **data-triggers** [**sap** *sap-id*] [**summary**]

Context show>service>dynsvc

Description This command displays the active dynamic services data trigger details.

Parameters *sap-id* — Specifies the dynamic services SAP for which the details are shown.

summary — Limits the output to one line per data trigger SAP.

Output The following output is an example of data trigger information.

Sample Output

```
# show service dynamic-services data-triggers
===============================================================================
Dynamic Services Data-triggers
===============================================================================
SAP                         : 1/1/4:1214.101
-------------------------------------------------------------------------------
Acct session-ID             : 144DFF0000009156A24138
MAC                         : 00:51:00:dd:01:01
IP                          :
State                       : sapCreated
-------------------------------------------------------------------------------
No. of Data-triggers: 1
===============================================================================
```

Table 176 describes data trigger fields.

Table 176 **Data Trigger Field Descriptions**

Output field	Description
Acct session-ID	The RADIUS accounting session ID for this dynamic services data trigger. This accounting session ID is used as an accounting multi-session ID in RADIUS accounting for the associated dynamic services. It can also be used as a key in CoA or Disconnect Messages to set up or terminate associated dynamic services.
MAC	The MAC address learned to set up this dynamic service data trigger. The MAC address is included in the Access-Request message for RADIUS authentication.
IP	The IPv4 or IPv6 address learned to set up this dynamic service data trigger. If the data trigger packet was not an IP packet, then this field is empty. When available, the IP address is included in the RADIUS authentication and accounting messages.
State	The current state of the dynamic service data trigger: Pending—(initial state) data trigger received and authentication started Accepted—(transient state) authentication succeeded; dynsvc script started but not yet completed sapCreated—(final state) corresponding dynamic services sap created

dynamic-services-policy

Syntax **dynamic-services-policy** [*policy-name*]

Context show>service>dynsvc

Description This command displays the dynamic services policy information.

Parameters *policy-name* — specifies for which dynamic services policy the information is requested.

Output The following output is an example of dynamic service policy information.

Sample Output

```
# show service dynamic-services dynamic-services-policy "dynsvc-policy-1"
===============================================================================
Dynamic Service Policies
===============================================================================
Dynamic Service Policy      : dynsvc-policy-1
-------------------------------------------------------------------------------
cli-user                    : (Not Specified)
description                 : Dynamic Services Policy 1
script-policy               : script-policy-5
```

```
sap-limit                  : 2000

Accounting instance 1
Stats type                 : volume-time
server policy              : aaa-server-policy-1
Update interval (minutes)  : 30
Update interval jitter     : 180s

Accounting instance 2
Stats type                 : time
server policy              : aaa-server-policy-2
Update interval (minutes)  : 0
Update interval jitter     : 10%
-------------------------------------------------------------------------------
No. of Services-policies: 1
===============================================================================
```

Table 177 describes the Dynamic Services policy fields.

Table 177 **Dynamic Services Policy Field Descriptions**

Output Field	Description
dynsrv-policy-name	The unique name of a dynamic services policy, up to 32 characters.
cli-user	The identifier name of the CLI user associated with this Dynamic Services policy.
script-policy	The identifier name of the script policy associated with this Dynamic Services policy.
sap-limit	The limit of the number of SAPs (Service Access Point) that can be created using this Dynamic Services policy.
Stats type	The value used to identify the type of accounting statistics gathered, either **volume-time** or **time**.
server policy	The identifier name of a RADIUS server policy to be used for accounting.
update interval (minutes)	The time interval between consecutive accounting updates when using this Dynamic Services policy.
update interval jitter	The amount of jitter to be applied on the update interval.
No. of Services-policies	The total number Dynamic Services policies.

capture-sap

Syntax **capture-sap** [**statistics**] *sap-id*

Context show>service>id>dynsvc

Description This command displays the status and statistics of a dynamic services capture. Statistics include counters for the number of data triggers received and data trigger drop reasons.

Parameters *sap-id* — Specifies the dynamic services capture SAP.

statistics — Displays data trigger receive and drop reason counters for this dynamic services capture SAP.

Output The following output is an example of capture SAP information.

Sample Output

```
# show service id 10 dynamic-services capture-sap 1/1/1:10.*
===============================================================================
Dynamic Services Capture SAP 1/1/1:10.*
===============================================================================
Dynamic services policy      : dyn-svc-1
Administrative state         : in-service
-------------------------------------------------------------------------------
Last Mgmt Change             : 01/20/2016 12:54:06
===============================================================================
# show service id 10 dynamic-services capture-sap 1/1/1:10.*  statistics
===============================================================================
Dynamic Services Capture SAP 1/1/1:10.* Statistics
===============================================================================
Data packets received by SAP                              : 0
-------------------------------------------------------------------------------
Drop Reason Counters
-------------------------------------------------------------------------------
No policy configured at capture SAP level                 : 0
No authentication configured in policy                    : 0
Data-trigger already exists                               : 0
Lockout is active                                         : 0
Reached data-trigger system limit                         : 0
No memory available                                       : 0
Unsuccessful authentication                               : 0
No data-trigger SAP-id in authentication                  : 0
Corresponding dynamic SAP is not created                  : 0
===============================================================================
```

Table 178 describes the Capture SAP policy fields.

Table 178 **Capture SAP Field Descriptions**

Counter	Description
Data packets received by SAP	The number of dynamic service data triggers received on the capture SAP that reached the CPM.

Table 178 **Capture SAP Field Descriptions (Continued)**

Counter	Description
No policy configured at capture SAP level	There is no dynamic-services-policy configured at the capture SAP. This is required to determine the authentication destination.
No authentication configured in policy	The authentication section in the specified in dynamic services policy is missing or incomplete.
Data-trigger already exists	A new data trigger frame is received for an existing data trigger that is authenticated, but the corresponding dynamic SAP is not yet created. The new data trigger packet is dropped.
Lockout is active	The data trigger for this managed SAP is currently in a lockout state due to previous authentication failures.
Reached data-trigger system limit	The maximum number of dynamic service data triggers supported on the system is reached. Additional data triggers are dropped.
No memory available	There is not enough system memory available to process the data trigger.
Unsuccessful authentication	The authentication for a data trigger on this capture SAP failed or timed out.
No data-trigger SAP-id in authentication	The dynamic services data trigger SAP ID is not provided in authentication. This is a mandatory parameter.
Corresponding dynamic SAP is not created	The data trigger successfully authenticated but the corresponding dynamic SAP was not created. This is typically caused by a dynamic services script error.

root-objects

Syntax **root-objects**

Context show>service>dynsvc

Description This command displays the root objects created by dynamic data services.

Output The following output is an example of dynamic services root object information.

Sample Output

```
# show service dynamic-services root-objects
===============================================================================
Dynamic Service Root Objects
===============================================================================
OID prefix                  : svcRowStatus
OID index                   : .100000
Snippet name                : vprn
Snippet instance            : VRF-1
Orphan time                 : N/A
-------------------------------------------------------------------------------
No. of Root Objects: 1
===============================================================================
```

Table 179 describes the Root Objects fields.

Table 179 **Root Objects Field Descriptions**

Output field	Description
OID prefix	The corresponding SNMP OID prefix for this root object.
OID index	The corresponding SNMP OID index for this root object.
Snippet name	The name of the python function that created this root object. The name is set to N/A when the root-object is orphaned.
Snippet instance	The instance for which the python function with “Snippet name” created this root object. If the snippet is a result from a dynamic reference, then the snippet instance is the reference-id string passed in the dyn.reference(). If the snippet is not the result from a dynamic reference, then the snippet instance is the dynamic data service SAP-ID. The instance is set to N/A when the root object is orphaned.
Orphan time	The timestamp when the root-object became orphaned (root-object not deleted when corresponding teardown function is called) or N/A if the root-object is not orphaned.

saps

Syntax **saps** [**control-session** *acct-session-id*] [**port** *port-id*] [**dynsvc-policy** *policy-name*] [**summary**] [**sap** *sap-id*] [**svc-id** *service-id*]

Context show>service>dynsvc

Description This command displays Dynamic Services SAPs (instances) details.

Parameters **summary** — Displays a summary view only.

acct-session-id — Specifies control session accounting session id

port-id — Specifies Ethernet port.

policy-name — Specifies dynamic services policy.

sap sap-id — Specifies dynamic services SAP id.

service-id — The service ID of the dynamic service.

Output The following output is an example of dynamic service SAP information.

Sample Output

```
# show service dynamic-services saps
===============================================================================
Dynamic Services SAP's
===============================================================================
SAP                          : 1/1/1:1.901
-------------------------------------------------------------------------------
Acct session-ID              : 242FFF000001AE512CE4B6
Acct session-ID control      : 242FFF000001AB512CE4B6
Service                      : [100000]
Dynamic Services policy      : dynsvc-policy-1
Number of scripts executed   : 1
Number of scripts w success  : 1
Last script action           : setup
Time of last script action   : 2013/02/26 16:37:10
Parameters of last action    : vprn_1={'t':('VRF-1',65000,1000,'cpe-int-1','192.
                             : 168.20.1/24','2001:db8:cafe::1/64',901,901,910,92
                             : 0,'192.168.20.0/24','192.168.20.2')}

Accounting instance 1
Status                       : enabled
Stats type                   : volume-time
Update interval (minutes)    : 30

Accounting instance 2
Status                       : enabled
Stats type                   : time
Update interval (minutes)    : 0

-------------------------------------------------------------------------------
No. of SAP's: 1
===============================================================================


# show service dynamic-services saps summary
===============================================================================
Dynamic Services SAP's summary
===============================================================================
SAP                           Acct-Session-ID        Acct-Session-ID-Ctrl
-------------------------------------------------------------------------------
1/1/1:1.901                   242FFF000001AE512CE4B6 242FFF000001AB512CE4B6
-------------------------------------------------------------------------------
No. of SAP's: 1
===============================================================================
```

Table 180 describes Dynamic Services SAP fields.

Table 180 **Dynamic Services SAP Field Descriptions**

Output field	Description
SAP	The dynamic service SAP ID.
Acct session-ID	The dynamic service accounting session ID.
Acct session-ID control	The control channel accounting session ID.
Service	The dynamic service ID.
Dynamic Services policy	The policy referenced to setup the dynamic service.
Number of scripts executed	The number of times the script was executed for this dynamic service (setup, modify, revert or teardown).
Number of scripts w success	The number of times the script was executed successfully for this dynamic service.
Last script action	The setup, modify, revert, teardown.
Time of last script action	The timestamp of the last script action.
Parameters of last action	The content of the Dynamic Services Script Parameters attribute corresponding with the last action.
Status	RADIUS accounting is enabled or disabled.
Stats type	The type of statistics reported in accounting.
Update interval (minutes)	The time, in minutes, between Accounting Interim Update messages.

script

Syntax **script**

Context show>service>dynsvc

Description This command enables the context to show dynamic services script information.

snippets

Syntax **snippets** [**detail**]
snippets name *snippet-name* [**instance** *snippet-instance*] [**detail**]

Context show>service>dynsvc>script

Description This command displays the dynamic services snippets information. A snippet is the name of a set of Python functions to create, modify or destroy configuration; for example: a snippet called 'vprn' in order to create, modify, or destroy VPRN services.

The CLI output generated by a single dynamic service python function call is a snippet instance.

Parameters **detail** — Displays dynamic services snippet information

snippet-name — Specifies the snippet name.

snippet-instance — *S*pecifies the snippet instance.

Output The following output show examples of dynamic service snippets.

Sample Output

```
# show service dynamic-services script snippets
===============================================================================
Dynamic Services Snippets
===============================================================================
Name                        Instance                      Ref-count  Dict-len
-------------------------------------------------------------------------------
vprn                        VRF-1                         1          75
vprn_1                      1/1/1:1.901                   0          190
-------------------------------------------------------------------------------
No. of Snippets: 2
===============================================================================


# show service dynamic-services script snippets detail
===============================================================================
Dynamic Service Snippets
===============================================================================
Snippet            : vprn:VRF-1
-------------------------------------------------------------------------------
reference-count    : 1
dictionary-length  : 75

Root-object
-----------
oid                : svcRowStatus.100000

Reserved-id
-----------
id                 : service-id:100000
-------------------------------------------------------------------------------
Snippet            : vprn_1:1/1/1:1.901
-------------------------------------------------------------------------------
reference-count    : 0
dictionary-length  : 190

Referenced-snippet
------------------
snippet            : vprn:VRF-1
-------------------------------------------------------------------------------
No. of Snippets: 2
```

```
================================================================================
```

Table 181 describes snippets fields.

Table 181 **Snippets Field Descriptions**

Output field	Description
Name	The snippet name.
Instance	The snippet instance.
Ref-count reference-count	The number of times this snippet is referenced by other snippets.
Dict-len dictionary-length	The length of the python dictionary stored for this snippet.
Root-object	The object created by a Dynamic Service of which the parent is not created by a Dynamic Service.

statistics

Syntax **statistics**

Context show>service>dynsvc>script

Description This command displays dynamic service script statistics. Only non-zero values are shown.

The script statistics can be cleared with the “clear service statistics dynamic-services” command.

Output The following output is an example of Dynamic Services script statistics information.

Sample Output

```
# show service dynamic-services script statistics
================================================================================
Dynamic Services Script Statistics
================================================================================
Description                                                          Counter
--------------------------------------------------------------------------------
python scripts with 0 retries due to timeout                          46
setup    - jobs launched                                              16
setup    - jobs handled                                               16
setup    - success                                                    13
setup    - syntax error                                               1
setup    - execution failed                                           2
teardown - jobs launched                                              15
teardown - jobs handled                                               15
```

```
teardown - success                                                          14
teardown - syntax error                                                     1
-------------------------------------------------------------------------------
No. of Script Statistics: 10
-------------------------------------------------------------------------------
Last Cleared Time: 02/26/2013 09:59:07
===============================================================================
```

Table 182 describes Dynamic Services script statistics fields.

Table 182

Output field	Description
Description	The name of the script.
Counter	The number of times the script executed.
python scripts with 0 retries due to timeout	The number of retries due to timeouts.
setup	The setup script action. jobs launched — The number of times the setup jobs launched. jobs handled — The number of times the setup jobs were handled. success — The number of times the setup jobs were successful. syntax error — The number of setup syntax errors. execution failed — The number of times the setup failed.
teardown	The teardown script action. jobs launched — The number of times the teardown jobs launched. jobs handled — The number of times the teardown jobs were handled. success — The number of times the teardown jobs were successful. syntax error — The number of teardown syntax errors.

summary

Syntax **summary**

Context show>service>dynsvc

Description This command displays the global configuration summary for dynamic services:

- Service range
- Timers

Output The following output displays dynamic service summary information.

Sample Output

```
# show service dynamic-services summary
===============================================================================
Dynamic Services Summary
===============================================================================
range start                 : 100000
range end                   : 200001
setup timeout Access Accept : 30
===============================================================================
```

Table 183 describes Dynamic Services summary fields.

Table 183 **Dynamic Services Summary Field Descriptions**

Output field	Description
range start	The start service ID to define the service range for dynamic services.
range end	The start and end service ID to define the service range for dynamic services.
setup timeout Access Accept	The timeout, in seconds, after which a script is canceled for setup actions received via RADIUS Access-Accept messages.

sap-using

Syntax **sap-using** [**msap**] [**dyn-script**] [**description**]

Context show>service

Description This command displays SAP information.

Parameters **dyn-script** — Displays dynamic service SAPs information.

Output The following output displays information about SAPs using Dynamic Scripts.

Sample Output

```
# show service sap-using dyn-script
===============================================================================
Service Access Points
```

```
===============================================================================
PortId                          SvcId      Ing.  Ing.    Egr.  Egr.   Adm  Opr
                                           QoS   Fltr    QoS   Fltr
-------------------------------------------------------------------------------
[1/1/1:1.901]                   [100000]   901   ip4     901   ip4    Up   Up
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
Number of Dynamic Service SAPs : 1, indicated by [<sap-id>] [<svc-id>]
-------------------------------------------------------------------------------
===============================================================================


# show service sap-using dyn-script description
===============================================================================
Service Access Points
===============================================================================
PortId                            SvcId        Adm  Opr  Description
-------------------------------------------------------------------------------
[1/1/1:1.901]                     [100000]     Up   Up   This is a dynamic
                                                         service SAP
-------------------------------------------------------------------------------
Number of SAPs : 1
-------------------------------------------------------------------------------
Number of Dynamic Service SAPs : 1, indicated by [<sap-id>] [<svc-id>]
-------------------------------------------------------------------------------
===============================================================================
```

Table 184 describes SAP using Dynamic Scripts fields.

Table 184 **SAPs Using Dynamic Scripts Fields**

Output field	Description
PortID	The Ethernet port.
SvcId	The service ID of the dynamic service.
Ing QoS	The ingress QoS policy ID.
Ingr Fltr	The ingress filter ID.
Egr Qos	The egress QoS policy ID.
Egr Fltr	The egress filter ID.
Adm	The administrative state of the SAP.
Opr	The operational state of the SAP.
Description	The text string describing the SAP.
Number of SAPs	The number of SAPs matching in the show criteria.
Number of Dynamic Services SAPs	The number of Dynamic Service SAPs matching in the show criteria.

15.5.2.4 Clear Commands

dynamic-services

Syntax **dynamic-services**

Context clear>service>stats

Description This command resets the dynamic services script statistics. See also **show service dynamic-services script statistics**.

data-trigger

Syntax **data-trigger sap** *sap-id*

Context clear>service>dynsvc

Description This command deletes all dynamic services associated with a dynamic services data trigger.

Parameters *sap-id* — Specifies the dynamic services data trigger SAP to delete.

15.5.2.5 Debug Commands

dynamic-services

Syntax [**no**] **dynamic-services**

Context debug

Description This command enables the context to configure dynamic services debugging.

data-triggers

Syntax [**no**] **data-triggers**

Context debug>dynsvc

Description This command enables the context to configure dynamic services data trigger capture SAP debugging.

The **no** form of this command removes all dynamic services data trigger capture SAP debug configurations.

capture-sap

Syntax **capture-sap** *sap-id* [**encap-val** *qtag*[*.qtag*]] [**mode** *mode*]
no capture-sap *sap-id*

Context debug>dynsvc>data-triggers

Description This command enables or disables the generation of dynamic services data trigger debug events, such as:

- data trigger received
- authentication
- data trigger SAP created
- dynamic service SAP created
- dropped data trigger with drop reason such as data trigger exists or lockout active.

Multiple capture SAPs can be specified simultaneously.

Optionally, a single **encap-val** per **capture-sap** can be specified to limit the output of the debug events to the data trigger events with the specified encapsulation.

Optionally, the debug output can be restricted to dropped data trigger events only.

Parameters *sap-id* — Specifies the dynamic services data trigger capture SAP for which debug events should be logged.

encap-val *qtag[.qtag]* — Optionally restrict the debug output to data trigger events with the specified encapsulation.

Values 1 to 4094

mode — Optionally restrict the debug output to specific events.

Values all—log all data trigger events
dropped-only—log only dropped data trigger events

scripts

Syntax **scripts**

Context debug>dynsvc

Description This command enables the context to configure dynamic services script debugging.

event

Syntax [**no**] **event**

Context debug>dynsvc>scripts

debug>dynsvc>scripts>inst
debug>dynsvc>scripts>script

Description This command enables/disables the generation of all dynamic data service script debugging events output: cli, errors, executed-cmd, warnings, state-change.

cli

Syntax [**no**] **cli**

Context debug>dynsvc>scripts>event
debug>dynsvc>scripts>inst>event
debug>dynsvc>scripts>script>event

Description This command enables/disables the generation of a specific dynamic data service script debugging event output: cli.

errors

Syntax [**no**] **errors**

Context debug>dynsvc>scripts>event
debug>dynsvc>scripts>inst>event
debug>dynsvc>scripts>script>event

Description This command enables/disables the generation of a specific dynamic data service script debugging event output: errors.

executed-cmd

Syntax [**no**] **executed-cmd**

Context debug>dynsvc>scripts>event
debug>dynsvc>scripts>inst>event
debug>dynsvc>scripts>script>event

Description This command enables/disables the generation of a specific dynamic data service script debugging event output: executed-cmd.

state-change

Syntax [**no**] **state-change**

Context debug>dynsvc>scripts>event
debug>dynsvc>scripts>inst>event

debug>dynsvc>scripts>script>event

Description This command enables/disables the generation of a specific dynamic data service script debugging event output: state-change.

warnings

Syntax [**no**] **warnings**

Context debug>dynsvc>scripts>event
debug>dynsvc>scripts>inst>event
debug>dynsvc>scripts>script>event

Description This command enables/disables the generation of a specific dynamic data service script debugging event output: warnings.

instance

Syntax **instance** *instance*

Context debug>dynsvc>scripts

Description This command enables the context to configure dynamic services script debugging for a specific instance.

Parameters *instance* — Specifies the instance name.

script

Syntax **script** *script*

Context debug>dynsvc>scripts

Description This command enables the context to configure dynamic services script debugging for a specific script.

Parameters *script* — Specifies the script name.

15.5.2.6 Tools Commands

dynamic-services

Syntax **dynamic-services**

Context tools>perform>service

Description This command enables the context to execute dynamic services tools perform commands.

evaluate-script

Syntax **evaluate-script sap** *sap-id* **control-session** *acct-session-id* **action** *script-action* [**dynsvc-policy** *name*] [**param-string** *string*]

Context tools>perform>service>dynamic-services

Description This tools command performs the execution of a dynamic service script action as if the corresponding RADIUS attributes were received from RADIUS. It is possible to setup, modify or teardown a dynamic service associated with the specified control channel.

Parameters *sap sap-id* — specifies the dynamic service SAP id.

control-session *acct-session-id* — Specifies the accounting session id of the control channel associated with this dynamic service

action *script-action* — Specifies the requested action: setup, modify or teardown.

dynsvc-policy *name* — Specifies the dynamic services policy to use for this action. Mandatory parameter for setup and modify actions. In case of a modify action, the dynamic services policy must be the same as the policy used at setup.

param-string *string* — Specifies the dynamic service parameter list. Mandatory parameter for setup and modify actions.

dynamic-services command-list

Syntax **dynamic-services command-list**

Context tools>dump>service

Description This command displays the list of supported commands that are allowed to be used in dynamic service CLI scripts.

There are two types of CLI nodes in this list:

- Pass through nodes: navigation is allowed but attributes creation or attribute changes are prohibited in this node.
- Allowed nodes: navigation, attribute creation and attribute changes are allowed in this node.

16 Diameter and Diameter Applications

16.1 Restrictions

Diameter-Based Restrictions:

- Accounting (RFC 6733, *Diameter Base Protocol*) via Diameter is not supported in this release.
- Accounting-Request (ACR), Accounting-Answer (ACA), Session-Termination-Requests (STR) and Session-Termination-Answer (STA) messages are not supported.
- SCTP and IPsec as transport protocols are not supported. TCP is supported.

Gx-Based Restrictions:

- Static hosts L2-Aware NAT hosts, PPPoA hosts, PPPoEoA hosts, MLPPP sessions and LAC (L2TP) hosts are not supported in Gx.
- Bridged homes and AA subscribers — Since there is no concept of a subscriber host in AA, the last AA policy submitted by Gx for any ESM subscriber host within the home is applied to the AA subscriber as a whole and overwrite any previously active AA policy.
- The <SAP, MAC> combination must be unique for each host (single stack or dual-stack).
- When an SLA profile instance contains multiple subscriber hosts, it is mandatory that all hosts have the same PCC rules applied.
- The Charging-Rule-Name within the Charging-Rule-Definition cannot contain a double colon (::) set of characters in the name string. The use of double colon in the name string itself is reserved for future use.
- Reporting about successful rule activation on the node (3GPP 29.212, §4.5.2) is not supported. The rule report is sent only if the rule instantiation fails.
- Time-based Usage-Monitoring is not supported.
- Gx persistency is not supported. However, upon reboot with ESM persistency enabled, the node will re-initiate Gx sessions (new CCR-I is generated for each Gx enabled host).
- Gy and Usage-Monitoring cannot be enabled for the same host and the same category-map at the same time. Gy is pre-configured at the time of the host instantiation. If a Usage-Monitoring request is received while Gy is enabled, the node will ignore the Usage-Monitoring request.

- Each ESM host can have up to sixteen Usage-Monitoring entities enabled simultaneously. Static configured categories and the internal category required for session level Usage-Monitoring are counted against this limit. The instantiation of the internal category for session level Usage-Monitoring is controlled with the **gx-session-level-usage** CLI command in the category map. If sixteen categories are configured, then Usage-Monitoring cannot be enabled per session (host) since this would exceed the limit of sixteen Usage-Monitoring entities per host.

16.2 Terminology

The term Gx interface (or simply Gx) is used to refer to the implementation of the Gx reference point on the node. Gx reference points are defined in the 3GPP 29.212 document.

The Enhanced Subscriber Management (ESM) subscriber is a host or a collection of hosts instantiated on the Broadband Network Gateway. The ESM subscriber represents a household or a business entity for which various services with committed Service Level Agreements (SLA) can be delivered.

AA Subscriber is a representation of ESM subscriber in MS-ISA for the purpose of managing its traffic based on applications (Layer 7 awareness). An AA subscriber has no concepts of ESM hosts.

BNG refers to the network element on which a Gx interface is implemented and policy rules are enforced (PCEF).

Flow – A flow in Gx context represents traffic matching criteria (traffic classification or traffic identification) based on any combination of the following fields:

- source IP address
- destination IP address
- source port or port ranges
- destination port or port ranges
- protocol field
- DSCP bits

A Gx flow is defined in the Flow-Information AVP:

```
Flow-Information ::= < AVP Header: 1058 >  3GPP 29.212 §5.3.53
          [ Flow-Description ]       3GPP 29.214 §5.3.8
          [ ToS-Traffic-Class ]      3GPP 29.212 §5.3.15
          [ Flow-Direction ]              3GPP 29.212 §5.3.65
          *[ AVP ]
```

Gx flows are similar to dynamically created filter or ip-criteria (QoS) entries and are inserted within the entry range configured for the base filter/qos-policy.

IP criterion – These fields are used in IPv4/v6 packet header used as a match criterion. The supported fields are DSCP bits and 5 tuple. This is part of traffic classification (or traffic identification) within the PCC rule or within the static qos-policy/filter entry.

Gx policy rule – There are three types of Gx policy rules supported:

- Gx based overrides — Subscriber related overrides (sub/sla/aa-profile, subscriber-id, QoS, filter, category-map, and so on).
- NAS filter entry inserts via Gx — Basic permit/deny entries, similar to ACL filter entries.
- PCC rules or IP-criterion based rules which are fully constructed Policy and Charging Control (PCC) rules with multiple QoS/filter actions per rule and its own traffic classification.

PCC rule represents a single or a set of IP based classifiers (DSCP bits or 5 tuple) associated with a single or multiple actions.

For example:

Each PCC rule can be removed via Gx from the node by referencing its name (Charging-Rule-Name AVP).

The PCC rule can contain a combination of QoS and IPv4/v6 filter actions as they pertain to the node.

PCC Rule Classifier — A flow-based (5 tuple) or a DSCP classifier defined in the Flow-Information AVP within the PCC Rule. There can be a single classifier or multiple classifiers (Flow-Information AVPs) within a single PCC rule. A PCC classifier (Flow-Information AVP) corresponds to an entry (match criteria) within the filter/ip-criteria definition.

CAM entry – A single entry in the CAM that counts toward the CAM scaling limit. For example a match condition within ip-criteria in a filter or qos-policy can evaluate into a single CAM entry or into multiple entries (where port-ranges are configured in the classifier, or where matching against UDP and TCP protocols are enabled simultaneously).

Subscriber Host Policy – A collection of PCC rules (classifiers and actions), Gx overrides, NAS filter inserts and statically configured rules (CLI defined QoS or filter) that are together applied to the subscriber host.

16.3 Diameter Base

An updated Diameter Base has been implemented to improve future enhancements. Advantages that the Diameter base implementation brings about are as follows:

- Future proof extensibility and better compliance with RFC 6733, *Diameter Base Protocol*
- CLI alignment with current AAA framework
- SMP (multi-core) capabilities
- Performance, the ability to process high volume of messages
- Diameter routing

The legacy Diameter base is still supported along with the new implementation, although the two cannot be used simultaneously for an application (NASREQ/Gx/Gy) within a given Diameter node instance within the SR OS. Operators are advised to deploy the new Diameter base as there are no plans going forward to add new functionality in the legacy Diameter implementation. Furthermore, the legacy Diameter implementation will be discontinued in one of the future SR OS releases, in accordance with the NOKIA software discontinuation policy.

This section describes the operation of the updated Diameter base. The principal difference between the two implementations from the operator's point of view is in the CLI. For this reason, the CLI descriptions are provided for both implementations. Some of the CLI commands remain the same for both implementations, and some of the commands are specific to their respective implementation (new or legacy). Each command that is specific to an implementation contains an explicit statement about the implementation to which it pertains.

16.3.1 Diameter Base Protocol

The Diameter base protocol is used to provide reliable and secure communication for Diameter applications between and across the Diameter peers (diameter clients, agents, or servers). Its routing capabilities to transport traffic across Diameter nodes rely on Diameter Identities which are comprised of host and realm names. Diameter applications supported in SR OS are NASREQ, Gx, and Gy.

In SR OS, Diameter base protocol runs over TCP. It starts by establishing a TCP connection with peers, followed by capability exchanges. A data exchange occurs in the form of Attribute Value Pairs (AVPs). Some of these AVPs are used by the base protocol itself, while others are used by Diameter applications that are layered on top of it.

16.3.2 Diameter Peers and the Role of a Diameter Node in SR OS

Diameter peers exchange messages via a transport level connection (TCP in SR OS). Each peer is identified by a unique Fully Qualified Domain Name (FQDN). On the IP level, Diameter peers are always directly connected.

Applications (NASREQ/Gx/Gy) relaying on Diameter base protocol use a concept of a session to communicate between the two end nodes that may be adjacent to each other or they may be multiple Diameter hops away.

This concept is shown in Figure 205.

Figure 205 **Diameter Peers**

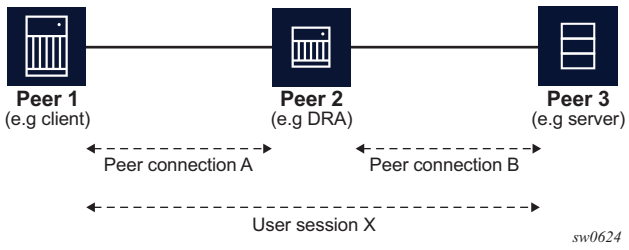

SR OS supports Diameter client functionality where it performs access control for a device running on the edge of the network. A Diameter client in SR OS can only initiate a connectivity request towards the peer but it will not accept a connectivity initiation request from the peer.

16.3.3 Capability Exchange

Upon TCP connectivity establishment, Diameter peers in SR OS exchange Capability Exchange messages. These messages contain information about the peer such as the peer's identity, supported applications, the protocol version, and security mechanisms. Capability Exchange messages do not traverse Diameter nodes but instead they are exchanged only between the peers (immediate Diameter next hops).

In this phase, the SR OS node performs peer authorization where the peer identity received from the peer using the Capability Exchange message is compared to the locally configured peer name. This peer identity (peer names) comparison in the SR OS is case insensitive. If the compared peer names do not match, the TCP session is closed.

The SR OS node always advertises all three supported applications (NASREQ, Gx, and Gy) in the Capability Exchange, regardless if these applications are configured in the SR OS node.

If no common applications are negotiated, the peer (the receiver of a Capability Exchange Request - CER) must return a Capability Exchange Answer (CEA) with the Result-Code AVP set to DIAMETER_NO_COMMON_APPLICATION and should disconnect the transport layer connection.

If the application ID of 0xffffffff (Diameter Relay Agent - DRA) is received by the SR OS, then this is interpreted as having a common application with the peer.

Diameter names and realms learned from peers through the Capability Exchange are used in the SR OS to populate local peer and realm routing tables that are used for forwarding and routing Diameter messages.

16.3.4 Connection Termination

The SR OS may initiate termination of the peering connection due to the following reasons:

- Peer shutdown — Graceful termination by initiating a DPR message from either side of the connection
- Watchdog timeouts — The peer is unresponsive and consequently the TCP connection is forcefully closed
- Message authorization failure during Capability Exchange phase with the peer — No support for common application or host names mismatch
- Failure to verify the following AVPs in incoming Diameter Base messages:
  - CEA: Origin-Host
  - DWR: Origin-Host, Origin-Realm, Origin-State-Id
  - DWA: Origin-Host, Origin-Realm, Origin-State-Id
  - DPR: Origin-Host, Origin-Realm
  - DPA: Origin-Host, Origin-Realm

  These values of these AVPs are checked against:
  - Origin-Host: peer name configuration

– Origin-Realm: learned from CEA
– Origin-State-Id: learned the first time seen in either CEA, DWR, DWA

16.3.5 Diameter Hosts and Realms

Diameter realms are administrative domains that have a relationship with the user account. For example, a Diameter client can be common to multiple departments with different domain or realm names, for example, xyz.com and wvu.com. In some cases, realms can represent different geographical regions. Regardless of the case, the realm names are maintained in Diameter nodes (as 'routes') and are used for routing of Diameter messages.

In general, realms can be regarded as a string in Network Access Identifier (NAI) that immediately follows '.' character (for example, jdoe.example.com). In SR OS, the realms are explicitly configured with a management interface as part of the Diameter configuration.

16.3.5.1 Forwarding/Routing of Diameter Messages

Forwarding and routing of the Diameter messages depends on the message type (requests vs answers). Diameter requests are forwarded/routed based on destination hosts/realms.

Forwarding and routing of request messages is based on the Diameter hosts or realms and is dependent on the two tables maintained in each node between the source and destination:

- peer table contains a list of active peers (adjacent nodes)
- a realm-based routing table contains mapping of realms to peers

  The realm-based routing table is in the SR OS and is created per Diameter node.

The peer to which the Diameter request message is sent may be a two step process. First, if the request message contains the Destination-Host AVP, the peer table is checked to see if there is a matching peer available. If so (the destination is directly connected), the request message is forwarded to it, without further checking the routing table for the next hop. This is called the forwarding phase. If the matching host is not present in the peer table (the destination is not directly connected), or the

Destination-Host AVP is not present in the request message (for example, CCR-I), then the Diameter routing table is checked for the matching realm. This is referred to as the routing phase. Each routable request message must contain the Destination-Realm AVP and the destination realm is a mandatory configuration parameter in SR OS for a given Diameter application within the SR OS Diameter node.

Diameter answer messages do not rely on the routing or forwarding tables. They are forwarded in the reverse direction of the matching requests in the transactional cache of each traversed Diameter node and the Hop-by-Hop AVP. Hop-by-Hop AVP is set to a locally unique value by each Diameter node that forwards request messages. This AVP is then used to match request and answer messages in the reverse direction. It allows a Diameter response to follow the same route as the corresponding Diameter request.

16.3.5.2 Configuration of Host Names and Realms in SR OS

Host names and realm names are essential in Diameter routing. They are conveyed in Diameter messages through the following AVPs:

- Origin-Host
- Origin-Realm
- Destination-Host
- Destination-Realm

Origin-Host and Realm AVPs are used to identify nodes in Diameter networks. Every Diameter message must carry these two AVPs. These two parameters are statically provisioned per SR OS node. An origin realm name is an optional CLI parameter, and if not configured it is extracted from the origin host name. The first “.” in the configured origin host name is used as a delimiter between the host portion and the realm portion. If there is no ‘.’ in the configured origin host name, then the origin host name is used as both, the origin host name and the origin realm name. Origin host and realm names do not support online changes. Once they are configured they cannot be changed without first deleting the Diameter node hierarchy. When received in Capability Exchange Answer messages, they are used to populate peer and realm tables in SR OS.

Destination-Host/Realm AVPs are used in forwarding or routing of Diameter request messages. The destination host name that is used by applications (Gx, Gy, and NASREQ) can only be learned from the incoming application level messages. The destination-realm is a mandatory configuration parameter in SR OS. The configured value is used in the initial request messages (CCR-I) when the destination host is not

yet known (has not been learned yet). However, the configured destination-realm can be overwritten for the session if that the value for the Destination-Realm AVP in an incoming application message differs from the configured value. In other words, the Destination-Realm AVP in an outgoing message is set to the value that is learned from the incoming message within the session.

Destination-Host or Destination-Realm AVPs are never present in:

- Diameter messages that are not to be proxied or relayed (such as CER)
- Diameter answer messages

16.3.5.2.1 Configuration

Origins are creation time parameters in the Diameter node CLI hierarchy. Multiple unique Diameter nodes can be instantiated within the same SR OS node.

```
configure>aaa>diam#
    node <origin-host-string> [origin-realm <origin-realm-string>]
```

Applications (Gx, Gy, and NASREQ) are associated with the Diameter node (and the origins) through the diameter-application-policy:

```
 configure>subscr-mgmt>diam-appl-plcy#
    diameter-node <origin-host-string> destination-realm <destination-realm-string>
```

The destination-realm referenced in the diameter-node above example is used to determine the next hop for the initial requests message. A peer with the matching realm must be configured and becomes the next hop for all messages destined to this destination realm.

16.3.5.2.2 Legacy Configuration

The centerpiece of legacy Diameter base configuration in SR OS is a diameter-peer-policy where all the Diameter base parameters related to communication with other external Diameter peers are defined (peering connections, DiameterIdentity, timers, and so on). This **diameter-peer-policy** is then associated with Diameter applications (NASREQ, Gx, and Gy) which relies on it to interact with the outside world.

In the legacy Diameter base, the origins are configured under the **diameter-peer-policy**:

```
configure>aaa>diam-peer-plcy#
    origin-host <origin-host-string>
    origin-realm <origin-realm-string>
```

while the destinations are configured one level below, under the peers:

```
config>aaa>diam>diameter-peer-plcy>peer#
    destination-host <destination-host-string>
    destination-realm <destination-realm-string>
```

If these parameters are not configured, the diameter-peer-policy is operationally down.

Applications (NASREQ, Gx, and Gy) are associated with the diameter-peer-policy through the diameter-application-policy:

```
configure>subscr-mgmt>diam-appl-plcy#
diameter-peer-policy <name>
```

16.3.6 Dynamically Learned Parameters

The Diameter configurations that are dynamically learned from a peer through the Capability Exchange are the received origin realm name (through the Origin-Realm AVP) and the App-ID. Both are used to populate the realm table (<realm, app-id> to peer mapping) and are used to find the next hop for the given destination-realm.

The Origin-Host AVP from the CEA is used to crosscheck the peer name in the peer table. Mismatches between the configured and learned (Origin-Host AVP in CEA) peer names causes TCP connection to be closed.

The destination-host name is learned by application in CCA-I and it can be re-learned through subsequent incoming application level messages. For example, the origin host name in incoming application level messages becomes the destination host name in outgoing application level messages.

16.3.7 Retransmissions and Message Timers

Message retransmissions are performed on the Diameter application level (NASREQ, Gx, and Gy) or on the TCP level. Diameter base itself does not retransmit messages.

Routable Diameter request messages can be retransmitted by the application (NASREQ, Gx, and Gy). If retransmission is enabled (by failover configuration on the application level), the original message is retransmitted only once and has the T-bit (retransmission bit) set.

An application level retransmission of the routable Diameter request messages may occur due to:

- Tx-timer timeout — An answer to a request message is not received within the time window determined by this configurable timer. This is the application level timer.
- An answer message is received with the E-bit set (protocol error)
- The path to the destination realm is unavailable

Answer messages are never retransmitted on the Diameter application level. However, they can be retransmitted on a TCP level.

Retransmission of non-routable Diameter request messages is driven by an internal timer which is set to a fixed value of 10 seconds. Non-routable Diameter messages are messages required for maintenance of the peering connection. They never propagate beyond the peer. This internal timer is used in the following cases:

- Failure to initiate TCP peering connection
- Uncomplete Capability Exchange — If the response to the CER message (CEA) is not received within 10 seconds, the peer connection is closed and the connection timer is started waiting for the new connection to be initiated
- Peer Disconnect Request — If the response to the DPR message is not received (DPA) within 10 seconds, the peer connection is closed

Diameter Watchdog messages receive special treatment. They are sent in absence of any traffic on the peer level and only when the peer level watchdog timer expires. However, only one watchdog message is sent, and it is not retransmitted. Watchdog messages are used to detect the liveliness of a peer.

Consecutive attempts to re-open a failed TCP peering connection are governed by the connection timer configured under the peer.

The following is a summary retransmissions and timers:

Retransmissions:

- are executed on TCP level for all Diameter messages using a TCP peering connection.
- Application level retransmissions apply only to routable Diameter request messages. Those retransmissions are governed by the application level Tx-timer.

Diameter watchdog message transmissions are governed by the peer level watchdog timer.

Timers include:

- a TX-timer for Application level routable Diameter request messages
- a watchdog timer for peer failure detection
- an internal timer for a DPR timeout, CER timeout, and TCP session establishment. In all three cases the peering session is closed.
- a connection timer for connection restart

16.3.8 Python Support

Python processing of the Diameter messages is supported in Diameter base where the processing can be enabled per message type per direction.

16.4 3GPP-Based Diameter Credit Control Application (DCCA) - Online Charging

The 3GPP-based Diameter Credit Control Online charging applications allow the control of subscriber access to services based on a pre-paid credit. The volume and time accounting on the node supports online charging using the Diameter Credit-Control Application (DCCA). The node supports Session Charging with Unit Reservation (SCUR), allowing the node to reserve volume and time quotas for rating-groups. Furthermore, the node supports centralized unit determination and centralized rating: it requests quotas and reports usage against the quota provided by the Online Charging Server (OCS). Credit control is always on a per-rating group basis. A rating group maps to a category inside a category map of the node volume-based and time-based accounting function.

The following are the basic configuration steps:

1. Configure a diameter policy

   In the **config>aaa** CLI context, configure a Diameter peer policy with one or multiple Diameter peers.

```
configure
    aaa
        diameter-peer-policy "diameter-peer-policy-1" create
            description "Diameter peer policy"
            applications gy
            connection-timer 5
            origin-host "bng.domain.com"
            origin-realm "domain.com"
            source-address 10.0.0.1
            peer "peer-1" create
                address 10.1.0.1
                destination-host "server.domain.com"
                destination-realm "domain.com"
                no shutdown
            exit
        exit
    exit
```

   When the diameter peer is reachable from IPv6, then the peer address should be specified as an IPv6 address. Optionally, an IPv6 source address can be specified:

```
configure
    aaa
        diameter-peer-policy "diameter-peer-policy-1" create
            ---snip---
            ipv6-source-address 2001:db8::1
            peer "peer-1" create
```

```
            address 2001:db8:100::1
            ---snip---
```

2. Configure a diameter application policy.

   In the **config>subscr-mgmt** CLI context, configure a diameter application policy:

   – Set the application to Gy (Diameter Credit Control Application),

   – Specify the Diameter peer policy to use and optionally specific additional Gy application specific parameters (for example AVP format).

```
configure
    subscriber-mgmt
        diameter-application-policy "diameter-gy-policy-1" create
            description "Diameter Gy policy"
            application gy
            diameter-peer-policy "diameter-peer-policy-1"
            gy
                avp-subscription-id subscriber-id type e164
                include-avp
                    radius-user-name
                exit
            exit
        exit
    exit
```

3. Create a category-map in which you define:

   – The credit type (time or volume).

   – A category defining the queues to monitor for quota consumption and the rating-group this category maps to in DCCA.

```
configure
    subscriber-mgmt
        category-map "cat-map-1" create
            description "Category Map"
            credit-type time
            category "cat-1" create
                rating-group 1
                queue 1 ingress-egress
                exhausted-credit-service-level
                    pir 256
                exit
            exit
        exit
    exit
```

4. Create a credit control policy.

   Define the credit control servers to use by specifying the diameter application policy. Optionally, specify the default-category-map and an out-of-credit-action.

```
configure
    subscriber-mgmt
        credit-control-policy "cc-policy-1" create
            description "Credit Control Policy"
            credit-control-server diameter "diameter-gy-policy-1"
            default-category-map "cat-map-1"
            out-of-credit-action change-service-level
        exit
    exit
```

5. Configure the diameter credit-control-policy in the sla-profile of the subscriber host for which credit control should be activated.

```
configure
    subscriber-mgmt
        sla-profile "sla-profile-3" create
            description "SLA profile"
            credit-control-policy "cc-policy-1"
        exit
    exit
```

The following are examples of Diameter online charging flows:

Scenario 1 — Depicts a redirect use-case:

When the quota is depleted, the subscriber is redirected to a web portal. When the credit is refilled, the OCS server notifies the BNG and provides a new quota. The configured out-of-credit-action when receiving a Final Unit Indication with action different from terminate is installed. See Figure 206 and Figure 207.

Figure 206 **Online Charging Scenario 1 - Redirect (1/2)**

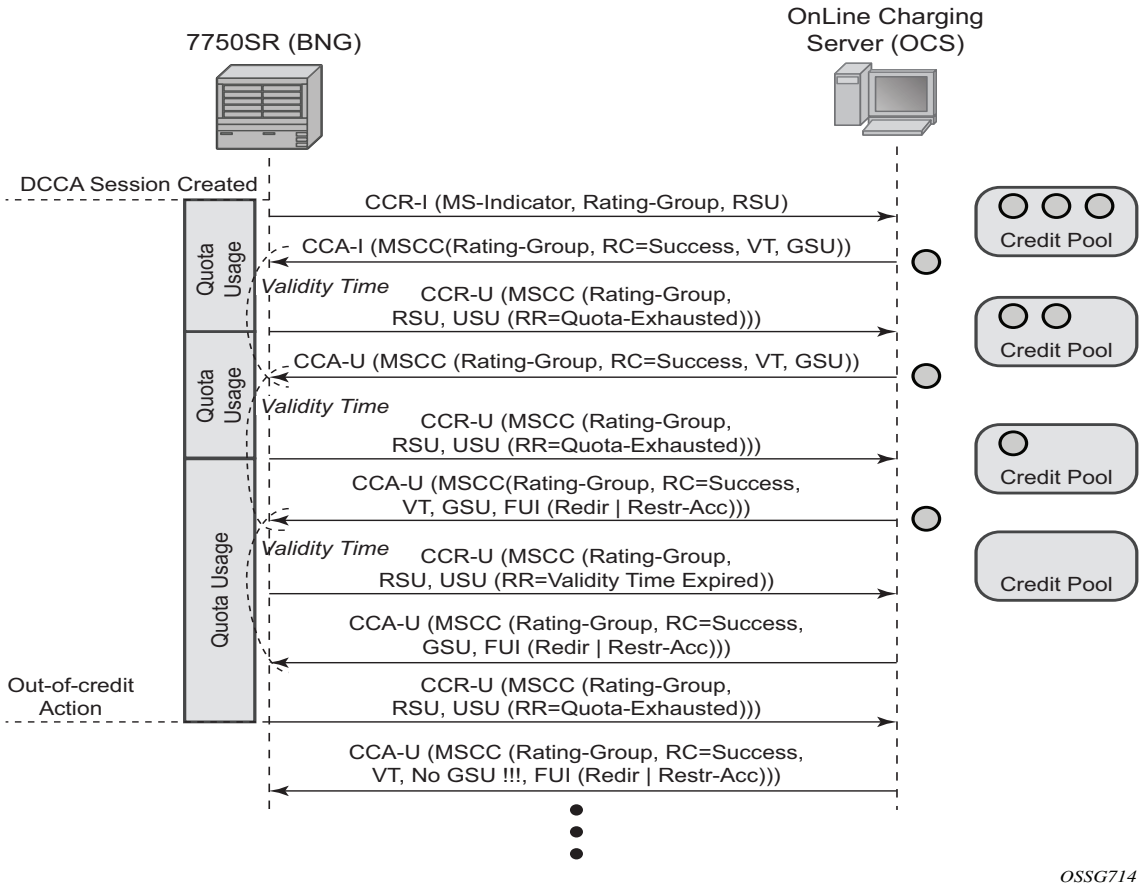

Figure 207 **Online Charging Scenario 1 - Redirect (2/2)**

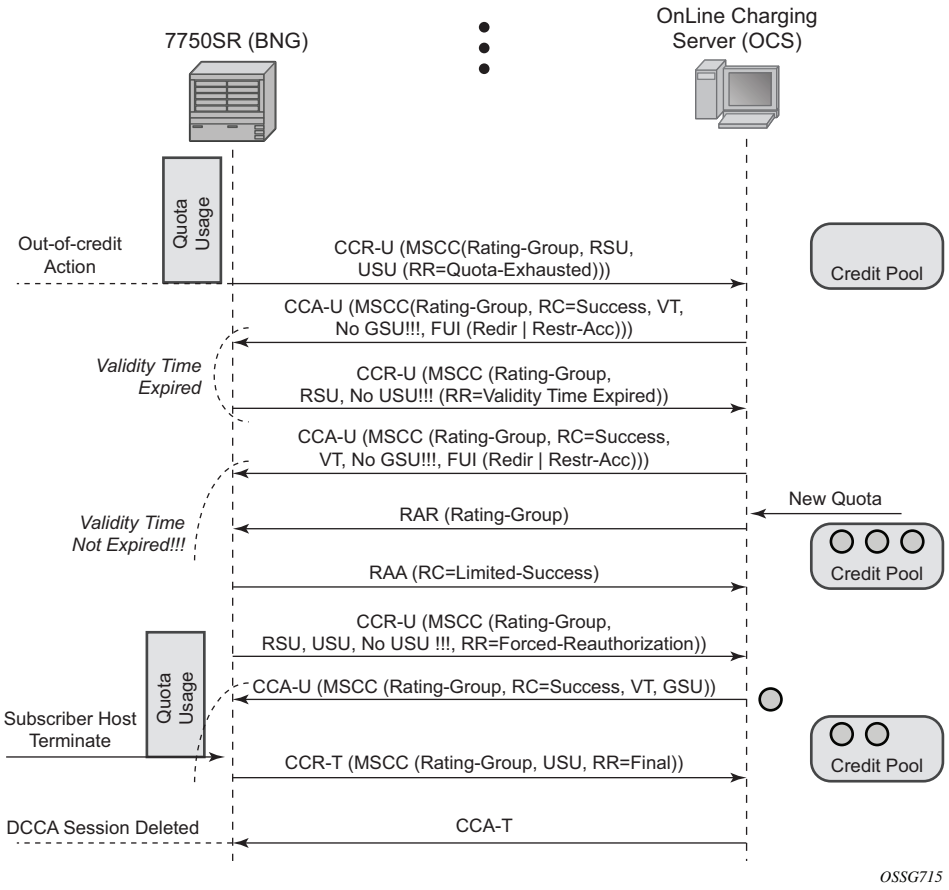

Scenario 2 — Depicts a terminate use case:

When the quota is depleted after reception of a Final Unit Indication with action set to Terminate, the subscriber host is disconnected. The configured out-of-credit-action is ignored in this case. See Figure 208.

Figure 208 **Online Charging Scenario 2 – Terminate**

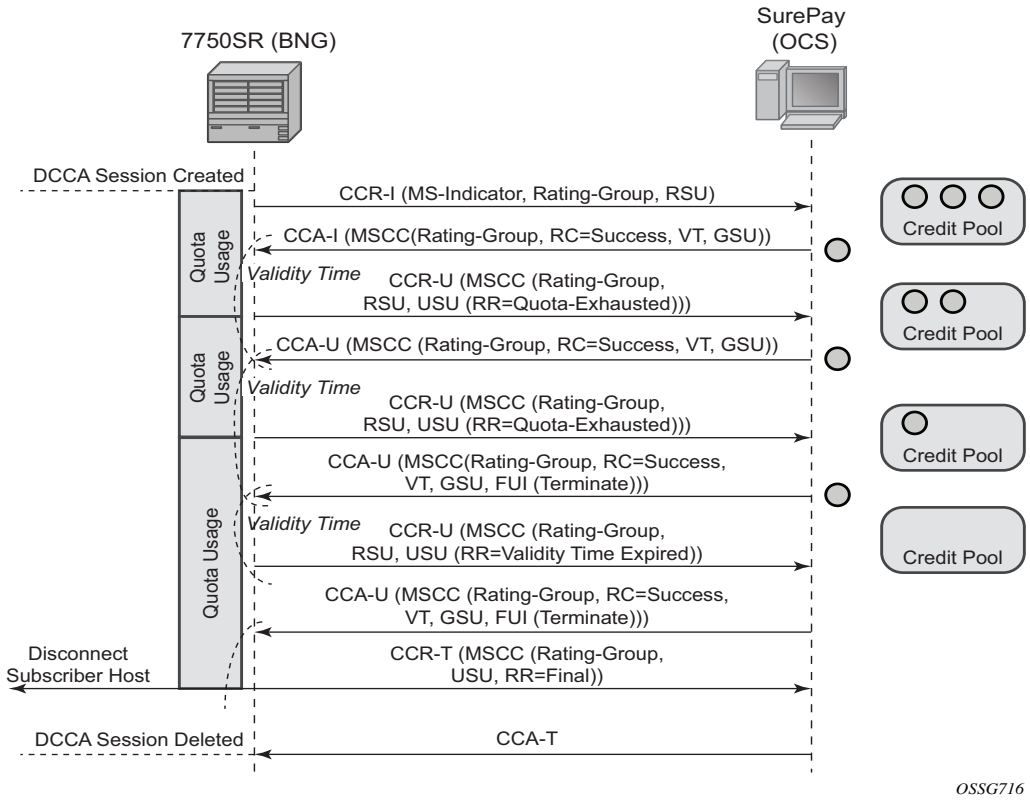

Abbreviations used in the previous drawings:

Table 185 **Online Charging Scenario Legend**

Abbreviation	Expansion
CCR	Credit Control Request (-Initial, -Update, -Terminate)
CCA	Credit Control Answer (-Initial, -Update, -Terminate)
RAR	Re-Authentication Request
RAA	Re-Authentication Answer
MSCC	Multiple Services Credit Control
GSU	Granted Service Unit
RSU	Requested Service Unit
USU	Used Service Unit
RC	Result Code

Table 185 **Online Charging Scenario Legend (Continued)**

Abbreviation	Expansion
RR	Reporting Reason
VT	Validity Time

16.4.1 Diameter Gy Out Of Credit Actions

When all quota of a Diameter Gy credit control rating group is consumed and no additional quota is granted by the server, the out of credit action as configured in the **credit-control-policy** or **category-map** is executed.

The out of credit action is one of the following:

- **continue**: traffic corresponding to the rating group is no longer accounted for by credit control
- **disconnect-host**: the subscriber host or session is disconnected if any one of the categories within the category-map has expired credit
- **block-category**: traffic corresponding to the rating group is blocked
- **change-service-level**: the service level is modified for the rating-group by applying one of the following:
  - a new PIR value to the corresponding queues or policers
  - an ingress or egress filter policy which could contain a filter to provide an HTTP redirect

16.4.1.1 Graceful Service Termination

RFC 4006, *Diameter Credit Control Application*, specifies a graceful service termination mechanism using the Final-Unit-Indication to indicate that an Answer message contains the final units for the service. When the final units are consumed, the action is specified with the Final-Unit-Action AVP.

```
Final-Unit-Indication ::= < AVP Header: 430 >
          { Final-Unit-Action }
         *[ Restriction-Filter-Rule ]
         *[ Filter-Id ]
          [ Redirect-Server ]
```

In SR OS, with the Final-Unit-Action AVP = TERMINATE, the subscriber host or session is disconnected. With the Final-Unit-Action = REDIRECT or RESTRICT_ACCESS, the **out-of-credit-action** as specified in the **credit-control-policy** or **category-map** is executed. In the case of REDIRECT, the URL specified in the Redirect-Server AVP is used when IPv4 HTTP redirect is enabled as the out of credit action and the following conditions are met:

- a Final-Unit-Indication AVP is present in the Multiple-Services-Credit-Control AVP of a CCA message
- the Final-Unit-Action AVP is set to REDIRECT (1)
- a Redirect-Server AVP is included with the following:
  - the Redirect-Address-Type AVP set to URL (2)
  - the Redirect-Server-Address AVP containing the URL to use for this rating group (**category-map**)
- the out of credit action for the corresponding rating group is set to **change-service-level** using one of the following CLI commands:
  - **configure subscriber-mgmt credit-control-policy** *policy-name* **out-of-credit-action change-service-level**
  - **configure subscriber-mgmt category-map** *category-map-name* **category** *category-name* **out-of-credit-action-override change-service-level**
- an IPv4 HTTP redirection action with **allow-override** is specified in the exhausted credit service level context for the corresponding rating group using the command **configure subscriber-mgmt category-map** *category-map-name* **category category-name exhausted-credit-service-level ingress-ip-filter-entries entry** *entry-id* **action http-redirect url allow-override**.

In all other cases, the URL specified in the Redirect-Server-Address AVP is ignored and the configured URL is used if HTTP redirect is enabled as the out of credit action.

The Restriction-Filter-Rule and Filter-Id AVPs included in the Final-Unit-Indication AVP are ignored.

Use the following **show** commands to find the active URL when an **out-of-credit-action change-service-level** is triggered that includes an IPv4 HTTP redirect action in the credit control ingress IP filter entries:

- **show service active-subscribers filter** [**subscriber** *sub-ident-string*]

  This command displays the active IPv4 ingress filter identifier.
- **show filter ip** *ip-filter-id*

With the filter identifier found in the previous command, this command displays the credit control inserted entries (Origin = “Inserted on ingress by Credit Control”), including the IPv4 HTTP redirect action with the static configured URL. The optional allow-radius-override flag may be shown. This flag is common for RADIUS and Diameter-based overrides, where the flag indicates that the URL may have been overridden by the Diameter Credit Control server. Use the next command to determine if an override was specified.

- **show service active-subscribers credit-control**

  This command displays the URL received from the Diameter Credit Control server for a rating group or category with “Out of Credit Action = ChangeServiceLevel” active. The URL is displayed in the “HTTP Rdr URL Override” field.

16.4.2 Extended Failure Handling (EFH)

In a Diameter Gy application, that is, a Diameter Credit Control Application (DCCA), Credit Control Failure Handling (CCFH) determines the behavior of a credit control client in fault situations. When the CCFH value is set to CONTINUE and a failure occurs, the credit control client first attempts a failover procedure. If failover is not enabled or not supported by both client and server, or the failover is not successful, the client deletes the credit control session and continues the service to the end user without the Diameter Gy session until the user disconnects.

Extended Failure Handling (EFH) enables the credit control client to establish a new Diameter Gy session with the Online Charging Server (OCS) in failure situations where CCFH is triggered and the CCFH value is set to CONTINUE.

The following occurs when EFH is enabled and the CCFH value is set to CONTINUE.

- Service to the end user continues during failures (such as lost connectivity to the OCS).
- A new Diameter Gy session is established when the failure is restored.
- Usage information is kept up-to-date for reporting purposes (optional).

Figure 209 shows an example of EFH.

Figure 209 **Extended Failure Handling (EFH)**

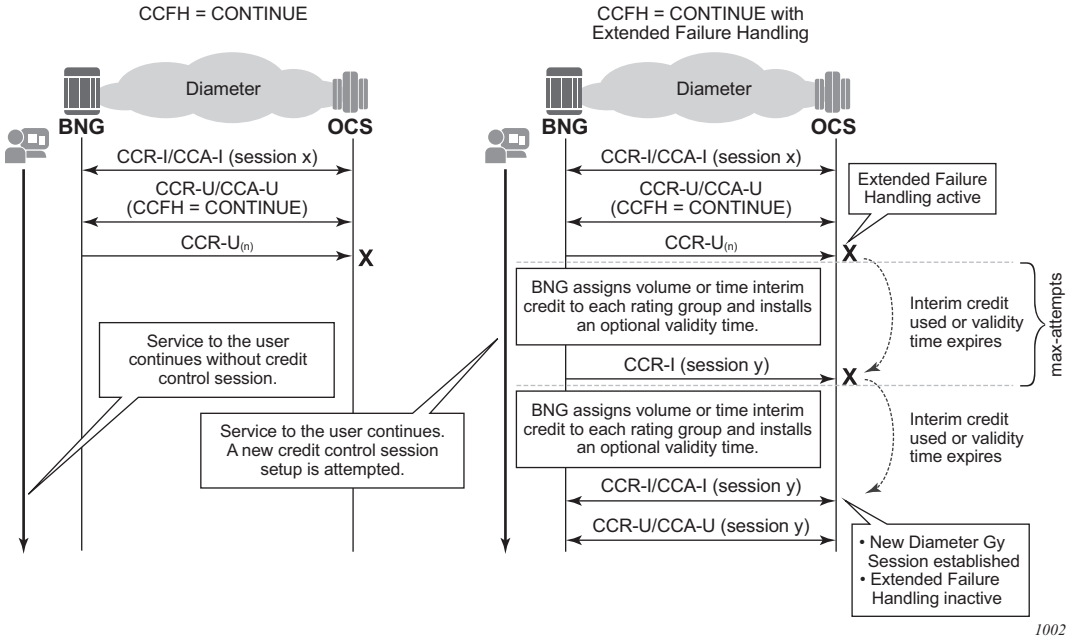

If a failure occurs when EFH is active, a preconfigured time interim credit or volume interim credit with an optional validity time is assigned to all rating groups. A new Diameter Gy session setup is attempted each time the interim credit is used or the validity time expires. The following occurs when an attempt to re-establish the Diameter Gy session is made.

- The user session continues normally and EFH becomes inactive when the Diameter Gy session is successfully established with the OCS.
- A new interim credit with optional validity time is assigned to all rating groups if the Diameter Gy session is not established with the OCS.
- The user session is terminated if the Diameter Gy session is still not established with the OCS after a configurable maximum number of attempts.

16.4.2.1 Extended Failure Handling Example Call Flow

Figure 210 **Extended Failure Handling - Example Call Flow**

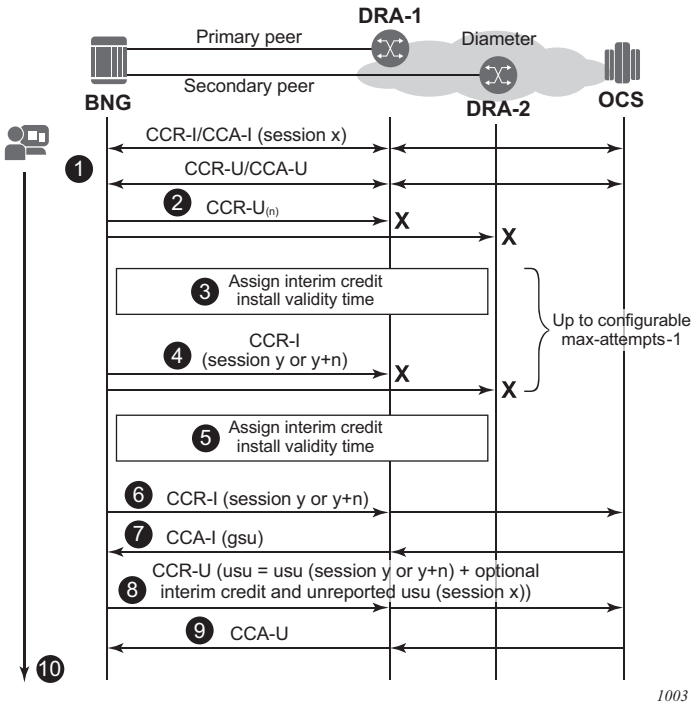

Figure 210 shows a sample call flow with EFH enabled. The following describes the call flow.

1. A Diameter Gy session is established between the Broadband Network Gateway (BNG), or credit control client, and the OCS, or credit control server.
2. A Credit Control Request Update (CCR-U) message is sent via the primary peer but no answer (CCA-U) is received. A timeout occurs which triggers a failover to the secondary peer. The same CCR-U is sent via the secondary peer. Again no answer is received. Because the CCFH value is set to CONTINUE for this session, EFH is activated.
3. The following EFH actions occur:
   - Service to the user continues uninterrupted.
   - If this is the first attempt to re-establish a Diameter Gy session:
     i. The failed Diameter Gy session (session ID x) is terminated without sending a CCR-T (Terminate) message.
     ii. The unreported used credits for each rating group are stored in an EFH unreported credit counter.

iii. A new Diameter Gy session (new session ID y) is created internally but is not yet established with the OCS; the CCR-I (Initial) message is sent later.

– For all subsequent attempts to re-establish a Diameter Gy session:

i. The failed Diameter Gy session (session ID y or y+n) is terminated without sending a CCR-T (Terminate) message.

ii. The unreported used interim credit for each rating group is added to the EFH unreported credit counter.

iii. A new Diameter Gy session is created internally but not yet established with the OCS: the CCR-I (Initial) message is sent later. The internal Diameter session is created with the same session ID (y) or a new session ID (y+n) based on a configuration knob.

– Pre-configured interim credit is assigned to all rating groups and an optional validity time is installed.

4. If either all interim credit is used or the validity time expires for one of the rating group, an attempt is made to establish the new Diameter Gy session (session ID y or y+n) with the OCS.

   A CCR-I message is sent via the primary peer but no answer (CCA-I) is received. A timeout occurs which triggers a failover to the secondary peer. The same CCR-I is sent via the secondary peer. Again no answer is received.

   EFH stays active for this user session.

   Steps 3 and 4 can be repeated multiple times until the maximum number of interim credit allocations is reached and the user session is terminated (not shown in this example call flow).

5. The EFH actions are as follows:

   – Service to the user continues uninterrupted.

   – The failed Diameter Gy session (session ID y or y+n) is terminated without sending a CCR-T (Terminate) message.

   – The unreported used interim credit for each rating group is added to the EFH unreported credit counter.

   – A new Diameter Gy session is created internally but is not yet established with the OCS: the CCR-I (Initial) message is sent later. The Diameter session is created with the same session ID (y) or a new session ID (y+n) based on the configuration.

   – Pre-configured interim credits are assigned to all rating groups and an optional validity time is installed.

6. If either all interim credits are used or the validity time expires for one of the rating groups, an attempt is made to establish the new Diameter Gy session (session ID y or y+n) with the OCS.

   A CCR-I message is sent via the primary peer.

7. An answer (CCA-I) is received with new granted service units (credit). Because communication with the OCS is restored, EFH becomes inactive.
8. The new Diameter Gy session resumes normal operation. Optionally, the EFH unreported credit usage is reported together with the usage from the newly granted credit in the next CCR-U credit negotiation for the rating group.
9. An answer (CCA-U) is received.
10. The service to the user continues and is uninterrupted during the OCS connectivity failure.

16.4.2.2 Extended Failure Handling Triggers

For EFH to become active, the Credit Control Failure Handling (CCFH) value must be set to CONTINUE.

For a new session, the CCFH value is set in the configuration:

```
configure
    subscriber-mgmt
        diameter-application-policy <application-policy-name> [create]
            on-failure [failover {enabled|disabled}] handling continue
```

For ongoing sessions, the CCFH value is determined from the configuration or can be overridden by the OCS by including the following AVP in an answer message (CCA-I or CCA-U):

```
[427] Credit-Control-Failure-Handling AVP = CONTINUE (1)
```

EFH is triggered when the CCFH value is set to CONTINUE and one of the following conditions occurs:

- a transmit failure—failure to send a CCR-I or CCR-U message

  Figure 211 shows an example of a transmit failure.

Figure 211 **EFH Trigger: Transmit Failure**

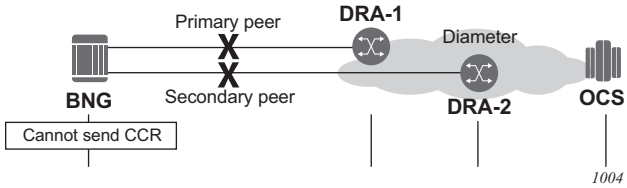

- a timeout—failure to receive an answer (CCA-I or CCA-U) within the configured timeout

Figure 212 shows an example of a timeout.

Figure 212 **EFH Trigger: Timeout**

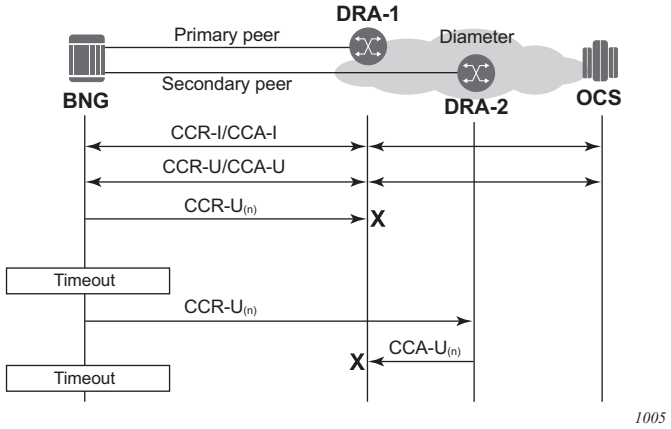

- a protocol error—failure due to a protocol error; seen as a Result-Code at command level in an answer message (CCA-I or CCA-U)

  Figure 213 shows an example of a protocol error.

Figure 213 **EFH Trigger: Protocol Error**

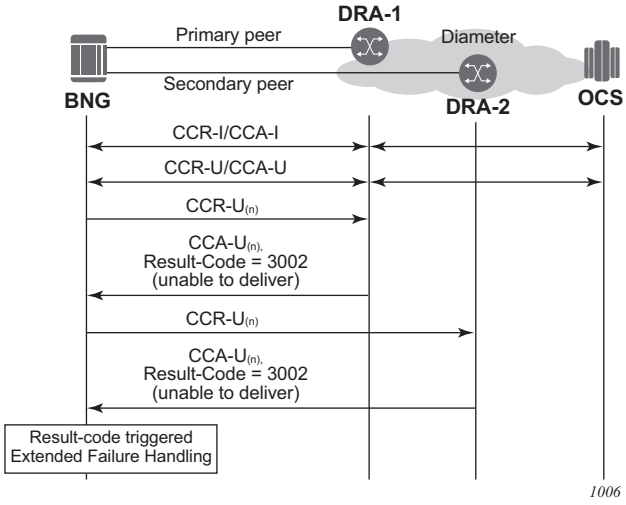

- a failure—reception of unknown Result-Code values in a Credit Control Answer message. Table 186 lists the known Transient and Permanent Failures Result-Code values.

Table 186 **Diameter Gy Known Transient and Permanent Failures Result-Code Values**

Result Code		Command Level		MCSCC Level
		CCA-I	CCA-U	
4001	DIAMETER_AUTHENTICATION_REJECTED	known	known	unknown
4010	DIAMETER_END_USER_SERVICE_DENIED	unknown	known	known
4011	DIAMETER_CREDIT_CONTROL_NOT_APPLICABLE	known	known	known
4012	DIAMETER_CREDIT_LIMIT_REACHED	unknown	known	known
5003	DIAMETER_AUTHORISATION_REJECTED	known	known	known
5030	DIAMETER_USER_UNKNOWN	known	known	unknown
5031	DIAMETER_RATING_FAILED	unknown	known	known

- a Diameter Gy message decoding error, for example (this is not an exhaustive list):
  - a missing Result-Code AVP
  - an unknown command code received
  - an incorrect session ID, origin host or origin realm in CCA
  - quota received for unexpected rating group
  - quota received with Result-Code 4012 (Credit Limit Reached)

16.4.2.3 Assigning Interim Credit

EFH interim credit can be specified in two ways:

- as a single volume interim credit value that is assigned to all rating groups of the Diameter Gy session with active EFH and that have no default credits configured

```
configure subscriber-mgmt
   diameter-application-policy <application-policy-name> create
      gy
         extended-failure-handling
            interim-credit
               volume <credits> {bytes|kilobytes|megabytes|gigabytes}
```

- as a single volume or time interim credit value per rating group (default credit)

  The activity threshold configured in the category map also applies to all rating groups that have time-based interim credit assigned while EFH is active.

```
configure subscriber-mgmt
```

```
   category-map <category-map-name> create
      activity-threshold <kilobits-per-second>
      category <category-name> create
         default-credit time <seconds>
         default-credit volume <credits> <bytes|kilobytes|megabytes|gigabytes>
```

A single validity time value can be specified and applied to all rating groups that have interim credit assigned, regardless of whether the interim credit is configured for all rating groups in the Diameter application policy or per rating group in the category map.

```
configure subscriber-mgmt
   diameter-application-policy <application-policy-name> create
      gy
         extended-failure-handling
            interim-credit
                 validity-time <seconds>
```

The maximum number of times that interim credit is assigned to all rating groups of a Diameter Gy session when EFH is active can be limited in the configuration. The **max-attempts** value also corresponds with the maximum number of attempts to establish a new Diameter Gy session with the OCS (the default maximum attempts = 10).

```
configure subscriber-mgmt
   diameter-application-policy <application-policy-name> create
      gy
         extended-failure-handling
            interim-credit
               max-attempts <count>
               max-attempts infinite
```

An attempt to establish a new Diameter Gy session with the OCS is made when one of the following conditions occurs.

- The assigned interim credit for one of the rating groups is used.
- The validity time of the interim credit for one of the rating groups expires.

When the maximum number of attempts is reached and a new Diameter Gy session is not yet successfully established, then the user session is terminated; that is, the corresponding subscriber hosts are deleted from the system.

The reporting of used EFH interim credit can be enabled using the following CLI command:

```
configure subscriber-mgmt
   diameter-application-policy <application-policy-name> create
      gy
         extended-failure-handling
            interim-credit
               [no] reporting
```

With reporting enabled, the following accumulated used credit is reported when a new Diameter Gy session is established with the OCS and usage reporting is triggered for a rating group:

- unreported used credit granted via the initial Diameter Gy session that caused the EFH activation
- used interim credit when EFH was active
- used credit granted via the new Diameter Gy session

By default, reporting is disabled and all the used credit from the initial Gy session and all used interim credit are discarded.

16.4.2.4 Enabling Extended Failure Handling

EFH is enabled by specifying **no shutdown** in the **extended-failure-handling** CLI context in a Gy Diameter application policy.

```
configure subscriber-mgmt
   diameter-application-policy <application-policy-name> create
      gy
         extended-failure-handling
            no shutdown
```

Table 187 lists the EFH states for a Diameter Gy session.

Table 187 **Extended Failure Handling States**

Extended Failure Handling State	Description
Disabled	EFH is disabled (**shutdown**). EFH cannot be triggered for the Diameter Gy session.
Enabled - active	EFH is enabled (**no shutdown**) and active. A failure event occurred that triggered EFH. Interim credit is assigned to all rating groups and when either the validity time expires or the interim credit is used for a rating group, a new attempt is made to establish a Diameter Gy session with the OCS.
Enabled - inactive	EFH is enabled (**no shutdown**) and inactive. A Diameter Gy session with the OCS is established. EFH is activated for the Diameter Gy session if a trigger condition occurs.

16.4.2.5 Configuration Example 1 - Single Volume Interim Credit Value

```
configure subscriber-mgmt
    diameter-application-policy gy-1 create
        --- snip ---
        on-failure handling continue
        gy
            --- snip ---
            extended-failure-handling
                no new-session-id          # default
                interim-credit
                    no reporting           # default
                    volume 100 megabytes
                    validity-time 900
                    max-attempts 96
                exit
                no shutdown
            exit
        exit
    exit
```

In this example, EFH is enabled and when active, 100 Mbytes of volume interim credit is assigned to all rating groups of the Diameter Gy session with a validity time of 900 s. The maximum number of attempts to establish a Diameter Gy session with the OCS is 96.

Therefore, a maximum of 96 x 100 Mbytes or 9.6 Gbytes can be consumed before the user session is terminated (that is, the subscriber host is deleted from the system) when the OCS remains unreachable. Alternatively, when less than 100 Mbytes per rating group is consumed every 15 minutes (900 s), the user is disconnected after 900 s x 96 = 24 hours when the OCS remains unreachable.

16.4.2.6 Configuration Example 2 - Interim Credit Values Per Rating Group

```
configure subscriber-mgmt
    category-map "cat-map-1" create
        activity-threshold 10
        category "cat-1" create
            description "Time interim credit per category"
            default-credit time 900
            rating-group 1
            queue 1 ingress-egress
        exit
        category "cat-2" create
            description "Volume interim credit per category"
            default-credit volume 10 megabytes
            rating-group 2
            queue 2 ingress-egress
```

```
        exit
        category "cat-3" create
            no default-credit
            rating-group 3
            queue 3 ingress-egress
        exit
    exit
    diameter-application-policy gy-1 create
        --- snip ---
        on-failure handling continue
        gy
            --- snip ---
            extended-failure-handling
                new-session-id
                interim-credit
                    reporting
                    volume 100 megabytes
                    validity-time 900
                    max-attempts 96
                exit
                no shutdown
            exit
        exit
    exit
```

In this example, EFH is enabled and when active, the following interim credit is assigned:

- 900 s of time interim credit to rating group 1 (category cat-1) with a validity time of 900 s. The activity threshold applies. When the usage stays below 10 kbytes/s, no time credit is used.
- 10 Mbytes of volume interim credit to rating group 2 (category cat-2) with a validity time of 900 s
- 100 Mbytes of volume interim credit to rating group 3 (category cat-3) with a validity time of 900 s

For inactive users, the validity time ensures that new Diameter Gy session connection attempts with the OCS are made at regular intervals.

For each attempt to establish a Diameter Gy session with the OCS, a new session ID is used.

When a new Diameter Gy session is successfully established, the EFH unreported credit for each rating group is included in the used service units when reporting is triggered for that rating group on the new Diameter Gy session.

When no new Diameter Gy session is established after 96 attempts, the user session is terminated; that is, the subscriber hosts are deleted from the system.

16.4.2.7 Monitoring the Extended Failure Handling State

Subscribers with Diameter Gy sessions that have EFH enabled can be displayed with the following CLI command:

```
# show service active-subscribers credit-control extended-failure-
handling [state] [summary]
```

where the EFH state can be set to:

- active—show subscribers with Diameter Gy sessions that have EFH enabled and EFH is active;
- inactive—show subscribers with Diameter Gy sessions that have EFH enabled and EFH is inactive; or
- all—show subscribers with Diameter Gy sessions that have EFH enabled and EFH is either active or inactive.

Example output:

```
A:BNG-1# show service active-subscribers credit-control extended-failure-handling
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber ipoe-msap-002 (sub-profile-1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:[1/1/4:1201.2] - sla:sla-profile-3
-------------------------------------------------------------------------------
Credit Control Policy: cc-policy-1
Category Map         : cat-map-1
Diameter Session Gy  : bng.domain.com;1464610029;840
CC Failure Handling  : continue
Extended Failure Handling (EFH)
  State              : active
  Attempts           : 1                      Maximum Attempts   : 10
  Active time        : 0d 00:00:10
  Total Active time  : 0d 00:00:10
  Total Active Count : 1
Category Name        : cat-1-time
Ingress Queues       : 1
Egress Queues        : 1
Ingress Policers     :
Egress Policers      :
Credit Volume Used   : 0                    Credit Time Used     : 11
Credit Volume Avail. : 0                    Credit Time Avail.   : 589
Credit Volume Thres. : 0                    Credit Time Thres.   : 0
Credit Expired       : False                Credit Negotiating   : False
Out Of Credit Action : None                 Quota Holding Time   : 0
Validity Time Used   : 0                    Validity Time Avail. : 0
EFH Unreported Credit
  Current Volume     : 0                      Current Time       : 601
  Total Volume       : 0                      Total Time         : 601
```

```
Category Name        : cat-2-volume
Ingress Queues       : 4
Egress Queues        : 4
Ingress Policers     :
Egress Policers      :
Credit Volume Used   : 38400               Credit Time Used     : 0
Credit Volume Avail. : 10447360            Credit Time Avail.   : 0
Credit Volume Thres. : 0                   Credit Time Thres.   : 0
Credit Expired       : False               Credit Negotiating   : False
Out Of Credit Action : None                Quota Holding Time   : 0
Validity Time Used   : 13                  Validity Time Avail. : 587
EFH Unreported Credit
  Current Volume     : 1527600               Current Time       : 0
  Total Volume       : 1527600               Total Time         : 0
-------------------------------------------------------------------------------
IP Address
                MAC Address          Session        Origin       Svc        Fwd
-------------------------------------------------------------------------------
10.1.1.101
                00:51:00:00:00:02    IPoE           DHCP         1000       Y
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
Number of active subscribers : 1
===============================================================================
```

The following information is displayed In the Extended Failure Handling (EFH) section of the example:

```
Extended Failure Handling (EFH)
  State              : active
```

State indicates that EFH is enabled and active. Another possible state is “inactive”. When EFH is disabled, no EFH information is included.

```
Attempts           : 1                     Maximum Attempts   : 10
```

Attempts indicates the number of times interim credit has been assigned to all categories followed by an attempt to establish a new Diameter Gy session with the OCS. When the attempt to establish a new Diameter Gy session with the OCS is still failing after the Maximum Attempts value is reached, then the user session is terminated (that is, the subscriber hosts are deleted from the system).

```
Active time        : 0d 00:00:10
```

Active time indicates the time since the EFH state became active for this subscriber session.

```
Total Active time  : 0d 00:00:10
  Total Active Count : 1
```

Total Active time indicates the accumulative time of all occurrences that EFH was active during the lifetime of this subscriber session.

Total Active Count indicates the number of times that EFH was active during the lifetime of this subscriber session.

For each category (rating group), the EFH Unreported Credit is displayed:

```
EFH Unreported Credit
  Current Volume     : 0                    Current Time       : 601
  Total Volume       : 0                    Total Time         : 601
```

The Current Volume and Current Time counters contain, respectively, the unreported volume and time credit for the current occurrence of the EFH in an active state. These counters include the unreported used credit for the initial Diameter Gy session that caused the EFH state to become active and the unreported interim credit from previous attempts. Used interim credit for the current attempt per category (rating group) is shown in the following counters:

```
Credit Volume Used   : 0                  Credit Time Used     : 11
Credit Volume Avail. : 0                  Credit Time Avail.   : 589
```

The Total Volume and Total Time counters contain respectively the accumulated total unreported volume and time credit for the previous occurrences of EFH active state. The total counters are updated when the EFH state toggles from active to inactive. When interim credit reporting is enabled, the counters are reset to zero when the actual usage reporting happens for that rating group. When interim credit reporting is disabled, the counters are accumulating the total unreported credit during the lifetime of the subscriber session.

Current and Total Volume EFH Unreported Credit counters are the sum of used ingress and egress octets.

For each category (rating group), the validity time is displayed:

```
Validity Time Used   : 13                 Validity Time Avail. : 587
```

The following fields are only displayed when the EFH state is active:

- Extended Failure Handling (EFH): “Attempts”, “Max Attempts” and “Active time”
- EFH Unreported Credit: “Current Volume” and “Current Time”

When EFH is disabled (shutdown), then the EFH information is not included in the credit control output.

16.4.2.8 Additional Call Flow Examples

16.4.2.8.1 User Disconnects While EFH is Active

The call flow in Figure 214 shows a scenario where EFH is activated before the session is established with the OCS. The scenario is similar when EFH is activated by a CCR-U message.

Figure 214 **EFH Call Flow - User Disconnects During EFH**

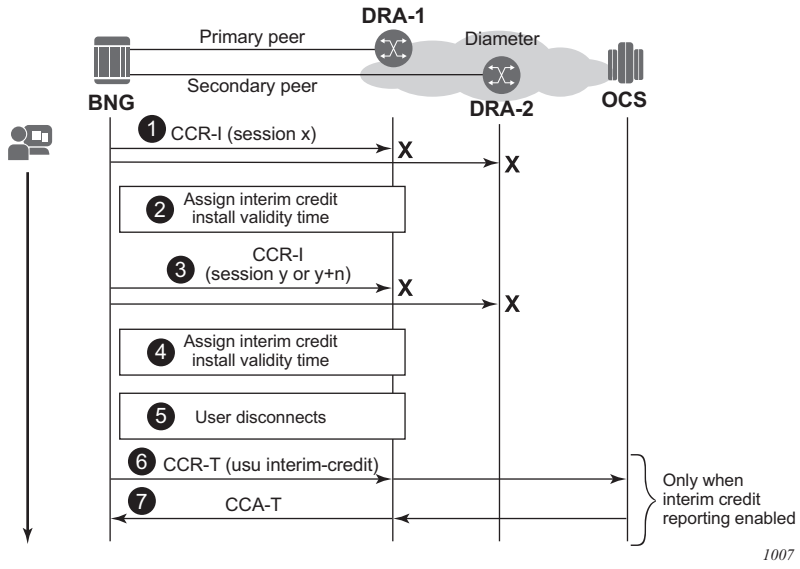

1. The initial Diameter Gy session setup fails: the CCR-I message sent to the primary peer times out or an error condition occurs that triggers Diameter Gy EFH. If failover is enabled, the CCR-I message is resent on the secondary peer. The CCR-I message sent to the secondary peer times out or an error condition occurs that triggers Diameter Gy EFH.
2. EFH becomes active for this user session. Interim credit is assigned to all rating groups and, optionally, a validity time is installed.
3. Interim credit is exhausted for a rating group or a validity time expires. A new attempt is made to establish the Diameter Gy session with the OCS; a CCR-I message is sent to the primary peer. A timeout or an error condition occurs that triggers Diameter Gy EFH to become active. If failover is enabled, the CCR-I message is resent on the secondary peer. The CCR-I message sent to the secondary peer times out or an error condition occurs that triggers EFH.
4. Interim credit is assigned to all rating groups and, optionally, a validity time is installed.

5. The user disconnects, resulting in a termination of the user session.
6. If interim credit reporting is enabled, a CCR-T is sent reporting the accumulated consumed interim credit.
7. If the OCS becomes reachable, a CCA-T may be received. Because the Diameter Gy session was not established with the OCS, the result code is DIAMETER_UNKNOWN_SESSION_ID.

16.4.2.8.2 The Maximum Number of Attempts is Reached

The call flow in Figure 215 shows a scenario where the maximum attempts is reached to establish a Diameter Gy session with the OCS.

Figure 215 **EFH Call Flow - Maximum Attempts Reached**

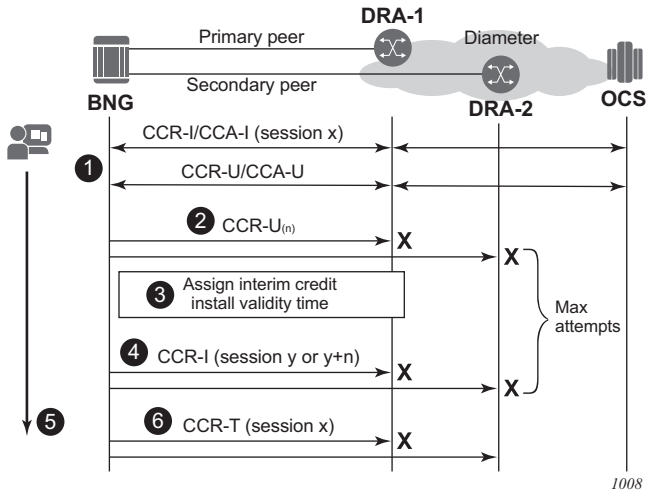

1. A Diameter Gy session is in progress.
2. An update message sent to the primary peer times out or an error condition occurs that triggers Diameter Gy EFH. If failover is enabled, the CCR-U message is resent on the secondary peer. The CCR-U message sent to the secondary peer times out or an error condition occurs that triggers Diameter Gy EFH.
3. EFH becomes active for this user session. Interim credit is assigned to all rating groups and, optionally, a validity time is installed.

4. Interim credit is exhausted for a rating group or a validity time expires. A new attempt is made to establish the Diameter Gy session with the OCS; a CCR-I message is sent to the primary peer. A new session ID is used for the first attempt. For subsequent attempts, a configuration command determines if a new or the same session ID is used. A timeout or an error condition occurs that triggers Diameter Gy EFH to become active. If failover is enabled, the CCR-I message is resent on the secondary peer. The CCR-I message sent to the secondary peer times out or an error condition occurs that triggers EFH.

   Steps 3 and 4 are repeated until the maximum attempts for interim credit is reached.
5. The user session is disconnected; that is, the subscriber hosts are deleted from the system.
6. When interim credit reporting is enabled, a CCR-T is sent, reporting the accumulated consumed interim credit: the unreported quota from the initial session x and the accumulated interim credit.

16.4.2.8.3 EFH Activation Triggered During Final Unit Indication (FUI)

This section describes two scenarios where EFH is activated during a graceful service termination initiated by the OCS with a FUI AVP. A graceful service termination with FUI action equal to REDIRECT or RESTRICT_ACCESS relies on a validity time or RAR to trigger a new credit negotiation with the OCS. Since the OCS is unreachable, it cannot be verified if a new quota has been granted. With EFH enabled, interim credit is assigned to guarantee service to the user until the connectivity with OCS is restored.

In the first scenario shown in Figure 216, the OCS initiates the graceful service termination with the Final-Unit-Action AVP = REDIRECT or RESTRICT_ACCESS. EFH is activated immediately after the out-of-credit action is installed.

Figure 216 **EFH Call Flow - Final Unit Indication Scenario 1**

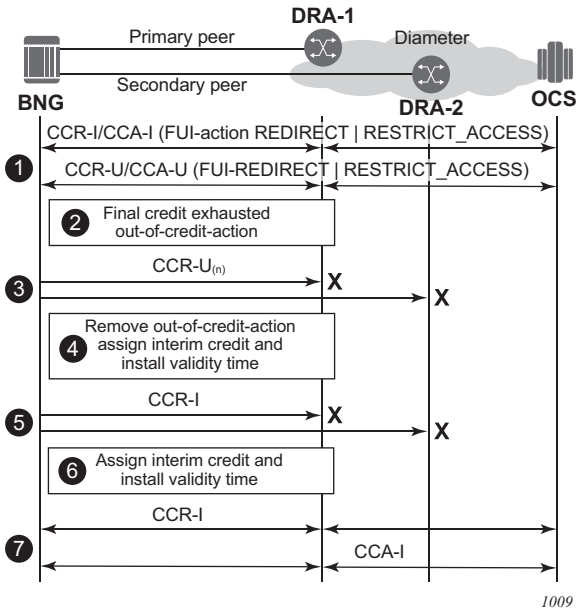

1. A Diameter Gy session is in progress. Either the CCA-I or a CCA-U contains the Final-Unit-Indication AVP and the Final-Unit-Action AVP is set to REDIRECT or RESTRICT_ACCESS.
2. The final assigned credit is exhausted and because the Final-Unit-Action AVP is different from TERMINATE, the provisioned out-of-credit action is installed. A CCR-U is sent to the OCS to notify it that the Final-Unit-Action has started.
3. The CCR-U message sent to the primary peer times out or an error condition occurs that triggers diameter Gy EFH. If failover is enabled, the CCR-U message is resent on the secondary peer. The CCR-U message sent to the secondary peer times out or an error condition occurs that triggers Diameter Gy EFH.
4. EFH becomes active for this user session. Interim credit is assigned to all rating groups and, optionally, a validity time is installed. The out-of-credit action is removed.
5. Interim credit is exhausted for a rating group or a validity time expires. A new attempt is made to establish the Diameter Gy session with the OCS and a CCR-I message (new session ID) is sent to the primary peer. A timeout or an error condition occurs that triggers Diameter Gy EFH to become active. If failover is enabled, the CCR-I message is resent on the secondary peer. The CCR-I message sent to the secondary peer times out or an error condition occurs that triggers EFH.
6. A new interim credit is assigned to all rating groups and, optionally, a validity time is installed.

7. Interim credit is exhausted for a rating group or a validity time expires. A new attempt is made to establish the Diameter Gy session with the OCS. A CCR-I is sent with the same or a new session ID to the primary peer. Whether the same or a new session ID is used is determined by a configuration command.

   A CCA-I answer message is received. EFH becomes inactive and the user session continues with the new established credit control session. Optionally, the used interim credit can be reported via the new established session.

In the second scenario shown in Figure 217, the OCS initiates the graceful service termination with the Final-Unit-Action AVP = REDIRECT or RESTRICT_ACCESS. EFH is activated after the FUI validity time expires.

Figure 217 **EFH Call Flow - Final Unit Indication Scenario 2**

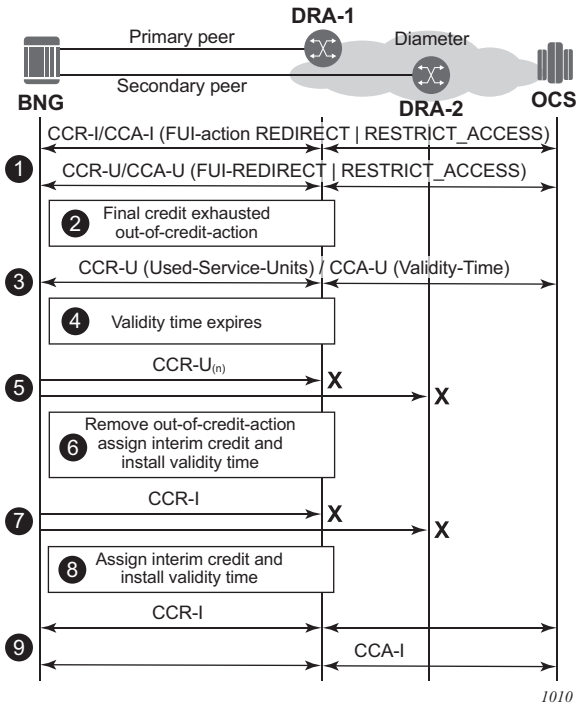

1. A diameter Gy session is in progress. Either the CCA-I or a CCA-U, contains the FUI AVP and Final-Unit-Action AVP = REDIRECT or RESTRICT_ACCESS.
2. The final assigned credit is exhausted and because the Final-Unit-Action is different from TERMINATE, the configured out-of-credit action is installed.
3. A CCR-U is sent to the OCS to notify that the Final-Unit-Action has started. The server responds with CCA-U containing a validity time.
4. The validity time expires.

5. The subsequent CCR-U message sent to the primary peer times out or an error condition occurs that triggers Diameter Gy EFH. If failover is enabled, the CCR-U message is resent on the secondary peer. The CCR-U message sent to the secondary peer times out or an error condition occurs that triggers Diameter Gy EFH.
6. EFH becomes active for this user session. Interim credit is assigned to all rating groups and, optionally, a validity time is installed. The out-of-credit action is removed.
7. Interim credit is exhausted for a rating group or a validity time expires. A new attempt is made to establish the Diameter Gy session with the OCS; a CCR-I message is sent to the primary peer. For the first attempt, a new session ID is used. For subsequent attempts, a configuration command determines if a new or the same session ID is used. A timeout or an error condition occurs that triggers Diameter Gy EFH to become active. If failover is enabled, the CCR-I message is resent on the secondary peer. The CCR-I message sent to the secondary peer times out or an error condition occurs that triggers EFH.
8. A new interim credit is assigned to all rating groups and, optionally, a validity time is installed.
9. Interim credit is exhausted for a rating group or a validity time expires. A new attempt is made to establish the Diameter Gy session with the OCS—a CCR-I message is sent to the primary peer.

   A CCA-I answer message is received. EFH becomes inactive and the user session continues with the new established credit control session. Optionally, the used interim credit can be reported via the new established session.

16.4.3 Gy CCR-t Replay

In diameter credit control, if the OCS is unreachable, EFH can be enabled to re-establish a Gy session with the OCS when the server becomes reachable again. This guarantees that no usage information is lost for billing. EFH applies to new and ongoing subscriber sessions.

If the OCS is unreachable after subscribers disconnect, the CCR-t that contains the final usage reporting is lost. CCR-t replay is a mechanism to recover final usage reporting data during OCS failure by replaying the CCR-t at configured intervals for a configured maximum lifetime or until an answer (CCA-t) is received from the OCS.

When enabled, CCR-t replay is triggered after:

- a CCR-t transmit failure occurs. If session failover is enabled, the CCR-t transmit is retried on the secondary peer.

- no CCA-t is received within the configured timeout. If session failover is enabled, the CCR-t is retried on the secondary peer.

Figure 218 illustrates CCR-t replay in action when the OCS is unreachable.

Figure 218 **CCR-t Replay**

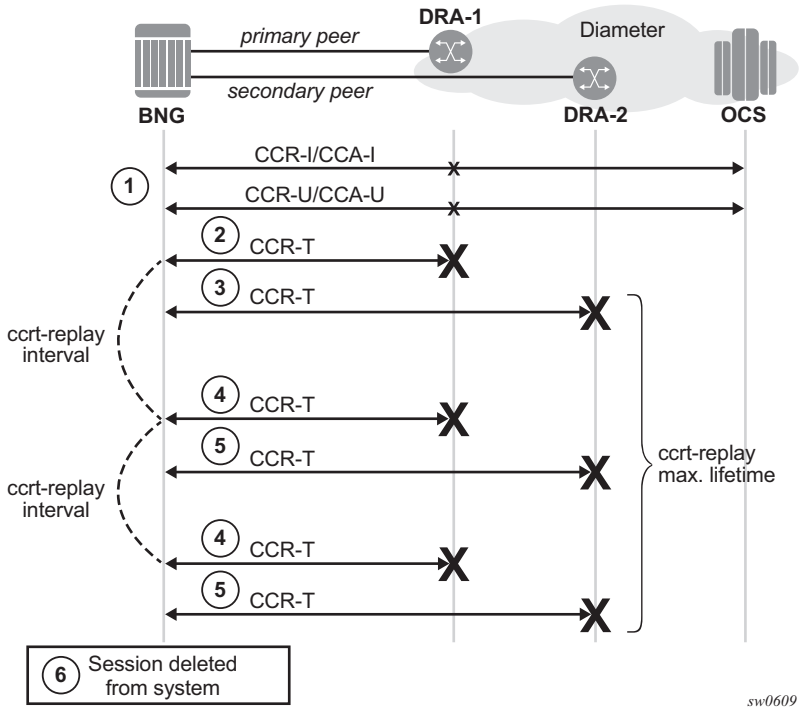

Where:

1. A diameter Gy session is in progress.
2. The subscriber disconnects, triggering the diameter Gy session to terminate. A CCR-t sent to the primary peer times out or an error condition occurs that triggers **ccrt-replay**.
3. If failover is enabled, the CCR-t is retried on the secondary peer. The CCR-t to the secondary peer times out or an error condition occurs that triggers **ccrt-replay**.
4. The CCR-t is replayed at the configured **ccrt-replay interval** until a valid answer is received or until the configured **ccrt-replay max-lifetime** is reached.
5. If failover is enabled, CCR-t replay is retried on the secondary peer.
6. The diameter Gy session in CCR-t replay is deleted when **crrt-replay max-lifetime** expires. An event is generated containing the Diameter Gx session ID.

The following configuration example enables CCR-t replay for Diameter Gy sessions:

```
configure subscriber-mgmt
    diameter-application-policy diam-gy-1 create
        gy
            ccrt-replay
                interval 1800
                max-lifetime 12
                no shutdown
            exit
        exit
    exit
```

The **ccrt-replay interval** value can be configured in seconds between 60 seconds (1 minute) and 86,400 seconds (24 hours).

The **ccrt-replay max-lifetime** value can be configured in hours between 1 hour and 24 hours.

Use the **show subscriber-mgmt diameter-session ccrt-replay** command to display Diameter Gy sessions that are in CCR-t replay mode.

Use the **show subscriber-mgmt diameter-session ccrt-replay diameter-application-policy** *name* command to display Diameter Gy sessions from the specified diameter application policy that are in CCR-t replay mode and per-diameter application policy statistics for those sessions.

Use the **clear subscriber-mgmt diameter-session ccrt-replay diameter-application-policy** *name* **sessions** command to drop all Diameter Gy sessions that are in CCR-t replay mode.

Use the **clear subscriber-mgmt diameter-session ccrt-replay diameter-application-policy** *name* **statistics** command to clear the CCR-t replay statistics and update the “Statistics last cleared time” timestamp.

16.5 Policy Management via Gx Interface

Gx is a reference point in the network architecture model describing mobile service delivery. The network elements are described in various technical documents under the umbrella of 3GPP and are used to deliver, manage, report on and charge end-user traffic for mobile users. The Gx reference point is used for policy control and charging control. As shown in Figure 219, it is placed between a policy server (Policy and Rule Charging Function (PCRF)) and a traffic forwarding node (Policy and Charging Enforcement Function) that enforces rules set by the policy server.

Figure 219 **Gx Reference Point**

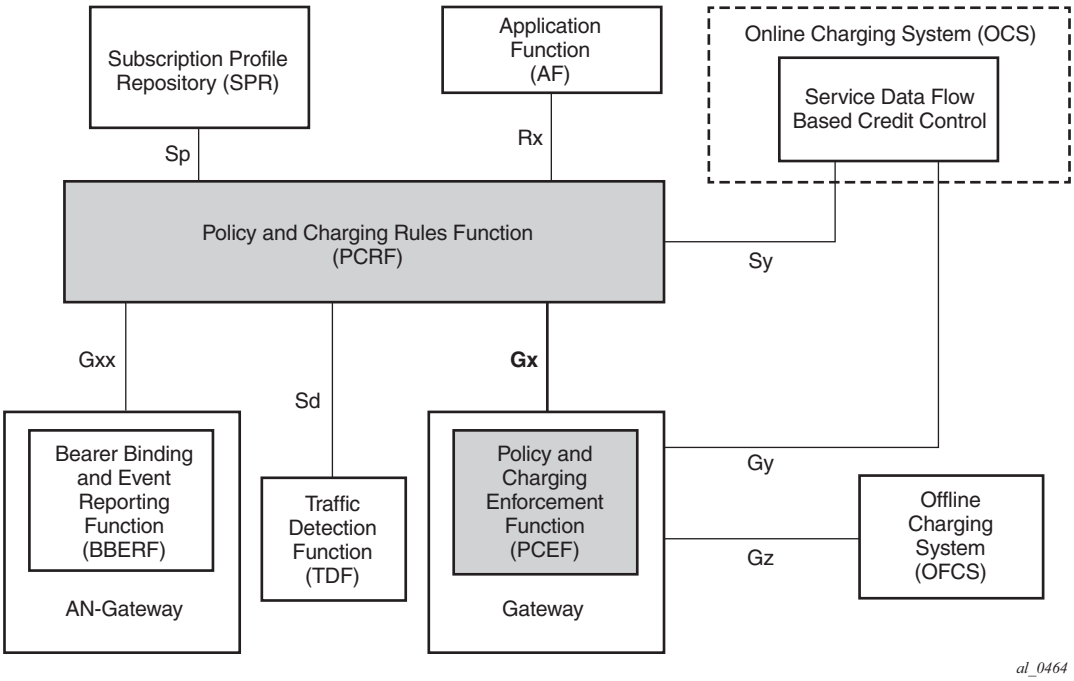

The Gx reference point is defined in the Policy and Charging Control (PCC) architecture within 3GPP standardization body. The PCC architecture is defined in the document 23.203 while the Gx functionality is defined in the document 29.212. Gx is an application of the Diameter protocol (RFC 3588/6733).

Although the Gx reference point is defined within 3GPP standardization body (spurred by mobile/wireless industry) its applicability has spread to wire-line operation as well. For example, mobile operators that also have fixed line customers (residential plus business) would like to streamline policy management in their mobile and non-mobile domains with a single and already existing Gx based policy management infrastructure. In other words they want to integrate policy management of nodes serving fixed line subscribers into the system that is currently managing mobile subscriber base.

In such mixed environments, the node plays the role of a PCEF with an integrated TDF (Traffic Detection Function, or Application Awareness [AA] in ALU terminology) where policies and charging rules can be managed via PCRF.

With WiFi Offload as a new emerging application, supporting Gx reference point on nodes is becoming even more important.

The Gx interface on the node encompasses the following functionality:

- Per subscriber host policy management
- Usage-Monitoring

Gx is applicable to Enhanced Subscriber Management (ESM) as well as to AA.

16.5.1 Gx Protocol

The Gx application is defined as a vendor specific Diameter application, where the vendor is 3GPP and the Application-ID for Gx application is 16777238. The vendor identifier assigned to 3GPP by IANA is 10415.

When a Diameter protocol defined over the Gx interface, the node (PCEF) acts as a Diameter Client and the PCRF acts as a Diameter server. The Gx Diameter application uses existing Diameter command codes from the Diameter Base Protocol (RFC 6733) and Diameter Credit Control Application (RFC 4006), both of which are already implemented on the node.

Gx is using Attribute-Value Pairs (AVPs) for data representation within its messaging structures (command codes). AVPs in Gx come from various sources:

- Gx-specific AVPs defined in 3GPP Gx specification TS 29.212.
- Reused AVPs from other existing Diameter applications (RFC 4006, RFC 4005, and so on), other 3GPP specifications, ETSI, and so on.
- RADIUS reused attributes (AVP codes 0-255 are reserved for RADIUS re-used attributes)

- Vendor-specific AVPs

The initialization and maintenance of the connection between the node (PCEF) and the PCRF is defined by the underlying Diameter protocol as defined in RFC 3588/ 6733.

16.5.2 Policy Assignment Models

Subscriber and AA policies on the node (PCEF with integrated TDF) is assigned through the Gx protocol from the policy server (PCRF).

There are two modes of operation:

- Pull mode — The policy creation or modification is solicited by the node.
- Push mode — The policy change is unsolicited by the node.

In the pull mode, during the host creation process, a user is authenticated by the AAA server. This process is independent from PCRF. Once the user is authenticated and the IP address is allocated to it, the node sends a request for policies to the PCRF via CCR-i messages (initialization request message). This communication occurs via the Gx interface. The subscriber-host must be uniquely identified in this request towards the PCRF. This sub- identification over the Gx interface could be by means of IP address, username, SAP ID, and so on.

Based on the user identification, PCRF submits policies to the node. Those policies can range from subscriber strings (sub/sla-profiles/AA-profiles) to QoS and filter-related parameters.

In the push mode, the PCRF initiates the mid-session policy change through the Re-Authentication Request (RAR) message (Figure 220).

Figure 220 **Policy Assignment Models**

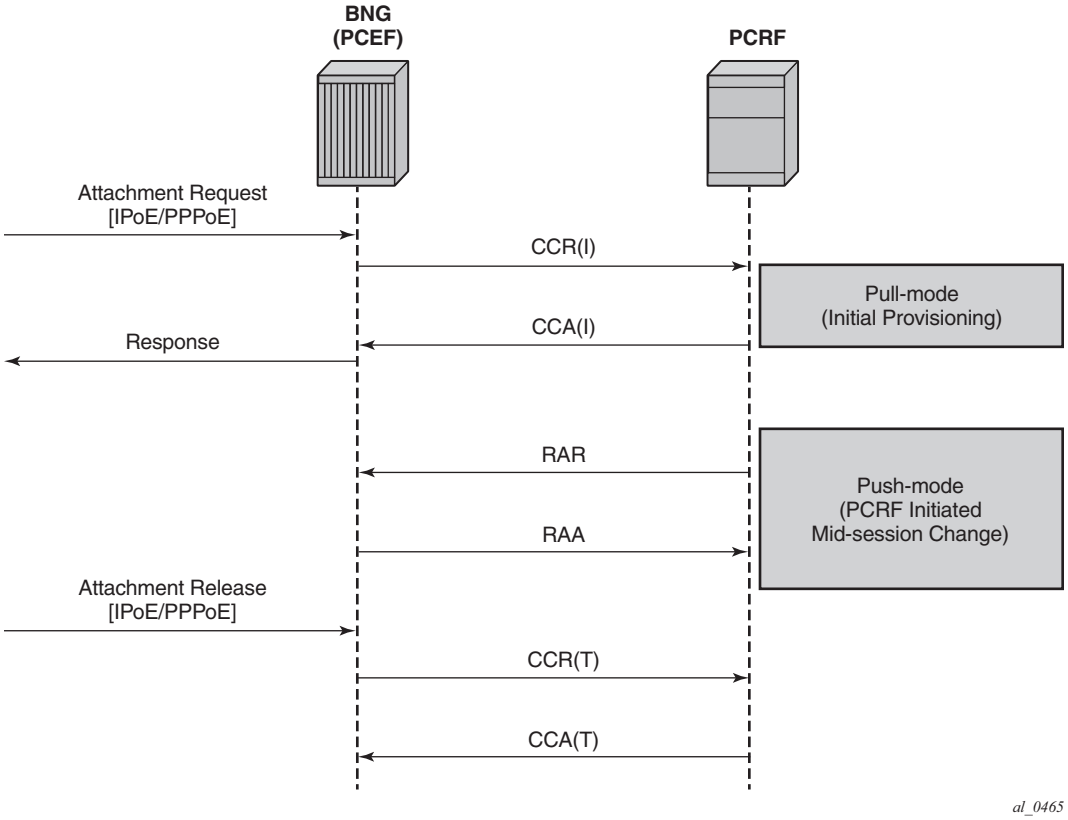

If that Usage-Monitoring is requested, the PCRF submitted policy changes are triggered by the Credit Control Request (Update) messages. This is based on ESM or AA Usage-Monitoring. Once the specified usage threshold is reached on the session-level, credit-category level, pcc rule level or application level on the node, the Usage-Monitoring is reported from the node to the PCRF in the CCR-u message. Refer to the *7450 ESS, 7750 SR, and VSR Multiservice Integrated Service Adapter Guide* for details on AA based Usage-Monitoring (Figure 221).

Alternatively, PCRF can request usage reporting on-demand via the **rar** command.

Figure 221 **On-Demand Usage Reporting**

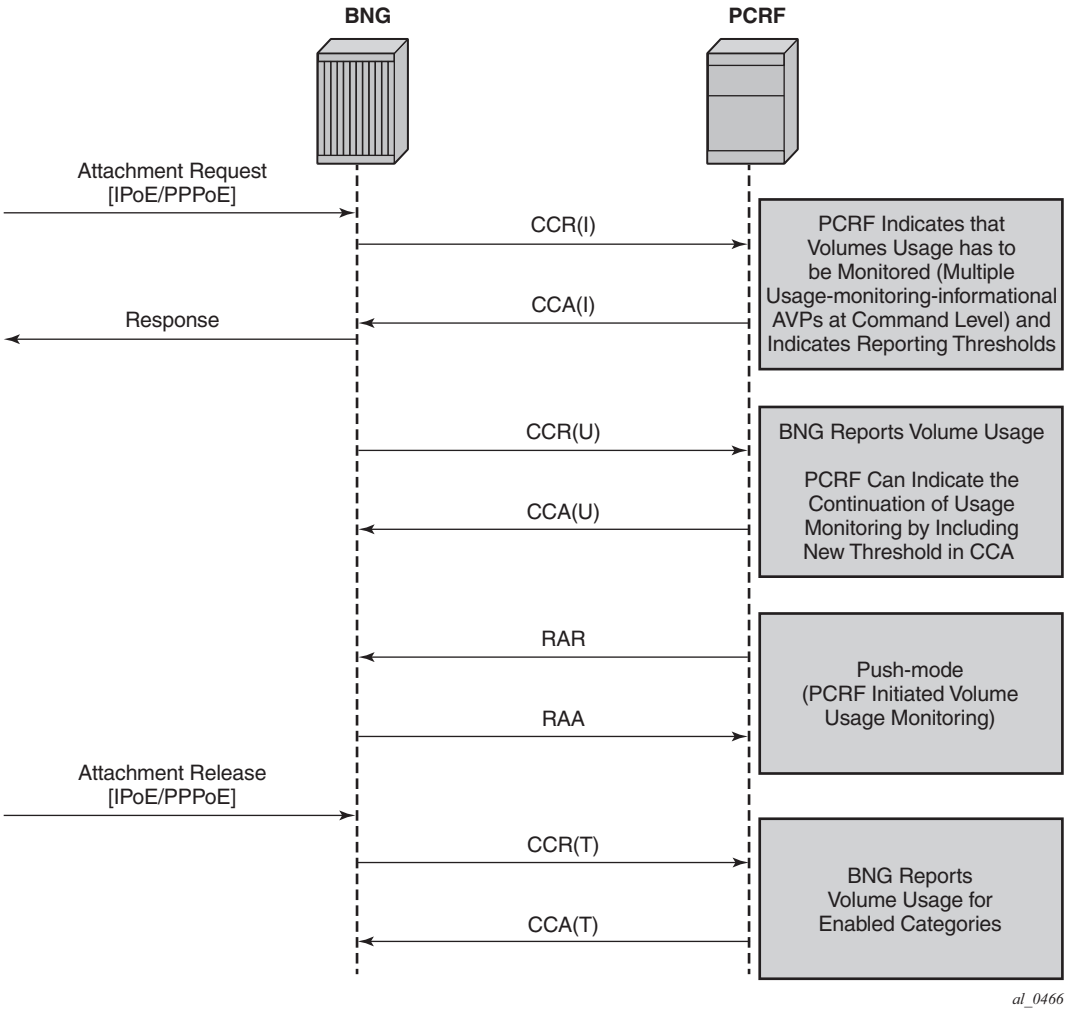

16.5.3 IP-CAN Session – Gx Session Identification

The IP Connectivity Access Network (IP-CAN) session is a concept that has roots in mobile applications. A policy rule via the Gx interface can be applied/modified to an entity that is identified as an IP-CAN session (in addition to individual bearers within the IP-CAN session, the bearer concept is currently not applicable to the BNG). For example, a UE (user interface or simply a mobile phone) can host several services, each of which appears as a separate IP-CAN session associated with the same IP address. For example, in mobile technologies, an IP-CAN session can be defined as <IP_address, APN, IMSI>, where:

- APN (Access Point Name) is the service identifier
- IMSI (International Mobile Subscriber Identification) is the UE identifier (SIM Card)

In wireline environment (ESM deployments), an IP-CAN session identifies an entity to which the policy can be applied or modified. In SR OS, this can be a single or dual stack IPoE host, IPoE session, or PPPoE session.

For the purpose of identifying the subscriber host or session in the SR OS node in all Gx-related transactions, the SR OS node generates a unique Gx session-id AVP (RFC 6733, §8.8) per single or dual stack IPoE host, IPoE session, or PPPoE session. Note that the Gx Session-Id AVP is not the same as the Acct-Session-Id attribute used in RADIUS accounting.

If the IPoE session is enabled, the Gx session key can be based on one of the following combinations (configuration dependent):

- {SAP,MAC}
- {SAP,MAC,Circuit-Id}; the Circuit-Id must be present in the packet
- {SAP,MAC,Remote-Id}; the Remote-Id must be present in the packet

If the IPoE session is disabled, the Gx session key for IPoE hosts (dual or single stack) is, by default, based on the {SAP,MAC} combination.

In an environment where a Layer 3 node is in front of the BNG, the MAC address of arriving packets are that of the Layer 3 node. Then, it is not possible to differentiate between IPoE hosts on the same SAP unless the IPoE session concept is enabled in SR OS. Each IPoE session must have a Circuit-id or Remote-id as a differentiator.

A session concept is native to PPPoE where the Gx session equates to a single or dual stack PPPoE session. This means that the Gx session key is based on a {MAC,SAP, PPPoE session-id} combination.

16.5.3.1 User Identification in PCRF

The following identification related AVPs are sent to the PCRF through Gx messages to aid in IP-CAN session identification:

- **subscription-id** AVP (RFC 4006, §8,46) — This can be used to identify the subscribers on the PCRF. For the supported fields within the subscription-id AVP, refer to the *7750 SR Gx AVPs Reference Guide*.
- **NAS-Port-Id** AVP (RFC 2869 / §5.17; RFC 4005 / §4.3)
- **AN-GW-Address** AVP (3GPP 29.212 / §5.3.49)

- **Calling-Station-ID** AVP (RFC 4005 / §4.6)
- **user-equipment-info** AVP (RFC 4006, §8.49)
- **logical-access-id** AVP (ETSI TS 283 034) — This contains the circuit-id from DHCPv4 Option (82,1) or interface-id from DHCPv6 option 18. The vendor-id is set to ETSI (13019).
- **physical-access-id** AVP (ETSI TS 283 034) — This will contain remote-id from DHCPv4 option (82,2) or DHCPv6 option 37. The vendor-id is set to ETSI (13019).

Physical and logical access IDs are also defined in BBF TR-134 (§7.1.4.1).

Table 188 shows PDP to PEP direction parameters.

Table 188 **PDP to PEP Direction Parameters**

Parameter	Category	Type	Description
Logical access ID	User identification	Octet String	The identity of the logical access to which the user device is connected. It is stored temporarily in the AAA function connected to PDP. This corresponds to the Agent ID in case of IPv4 and to THR Interface Id of DHCP option 82 for IPv6
Physical Access ID	User identification	UTF8String	The identity of the physical access to which the user device is connected. It is stored temporarily in the AAA function connected to the PDP. This corresponds to the Agent Remote ID

A Subscription-id AVP is most commonly used to identify the subscriber but any combination of the above listed parameters can be used to uniquely identify the IP-CAN session on PCRF and consequently identify the subscriber.

In addition, NAS-Port, NAS-Port-Type, and Called-Station-ID AVPs from RFC 4005 (§4.2, §4.4, §4.5) can be passed to the PCRF.

16.5.3.2 NAS-Port-Id as Subscription-Id

The node allows the NAS-Port-Id to be carried within the Subscription-Id AVP. Since the NAS-Port-Id may not be unique network-wide (two independent nodes may use the same NAS-Port-Id), there is a need for another identifier in conjunction with NAS-Port-Id to make the Subscription-Id unique across the network. This additional identifier is a custom string that can be appended to the NAS-Port-Id. It is defined when the NAS-Port-Id is configured for inclusion in Gx messages. Refer to the *7750 SR and VSR RADIUS Attributes Reference Guide* for information to format NAS-Port-Id AVP on the node.

The string can be used to identify the location of the node. For example:

```
@ALU-MOV-SITE-1
```

An example of the augmented NAS-Port-Id would look like:

```
NAS-Port-Id = lag-1.1/1/2:23.2000@ALU-MOV-SITE-1
```

where: 'lag-1.1/1/2:23.2000' is the part referencing the SAP on the node (port + vlan tags) and the '*@ALU-MOV-SITE-1*' is the node itself.

The NAS-Port-Id can be then inserted in the Subscription-Id AVP.

16.5.4 Gx Interface and ESM Subscriber Instantiation

Policy management via Gx enables operators to consolidate policy management systems used in wireline (mostly based on RADIUS/CoA) and wireless environment (PCRF/Gx) into a single system (PCRF/Gx).

The model for policy instantiation/modification via Gx is similar to the one using RADIUS CoA. The authentication and IP address assignment is provided outside of Gx while the policy management function is provided via Gx.

Some PCRFs may require that the IP address information is passed from the node in CCR-i. This assumes that the IP address assignment phase (via LUDB, RADIUS or DHCP Server) is completed before the PCRF is contacted via CCR-i. Message flow for various protocols (DHCP, AAA, Gx) related to IPv4 subscriber-host instantiation phase is shown in Figure 222.

CCR-i message is sent to the PCRF once DHCP Ack is received from the DHCP server. Relaying DHCP Ack to the client in the final phase of the host instantiation process depends on the answer from the PCRF and the configuration settings of the fallback function if the answer is not received.

This model allows the IP address of the host to be sent in the CCR-i message, even though the subscriber-host is not fully instantiated at the time when the CCR-i message is generated.

AAA/LUDB must still be used for authentication and assignment of parameters necessary to place the subscriber host in the proper routing context (service-id, grp-id, msap-policy).

Start of the accounting process nicely fits into this model since the host is not instantiated until the policy information from all sources (Gx, AAA, defaults) is known. Once the final sub-profile (containing the acct-policy) is known, the host is instantiated and accounting can consequently be activated.

The IP address itself cannot be assigned via Gx, and this functionality is outside of the Gx scope (3GPP TS 23.203 Rel12, Annex S, IP-CAN Session Establishment section).

The purpose of the CCR-i message is the following:

- To notify the PCRF that the sub-host was about to be instantiated on the node. Consequently, the PCRF creates a Gx session for the subscriber host if the CCR-i is successfully processed by PCRF.
- To identify the subscriber host in the PCRF, the PCRF will use the subscriber host identification information to identify the policy (for the subscriber host) that needs to be submitted to the node. The policy rules can be sent via CCA-i immediately following the initial CCR-i or they can be modified at any time during the subscriber-host lifetime through RAR messages.

Figure 222 shows the message flow during the DHCPv4 host instantiation phase.

Figure 222 **Messages Flow During DHCPv4 Host Instantiation Phase**

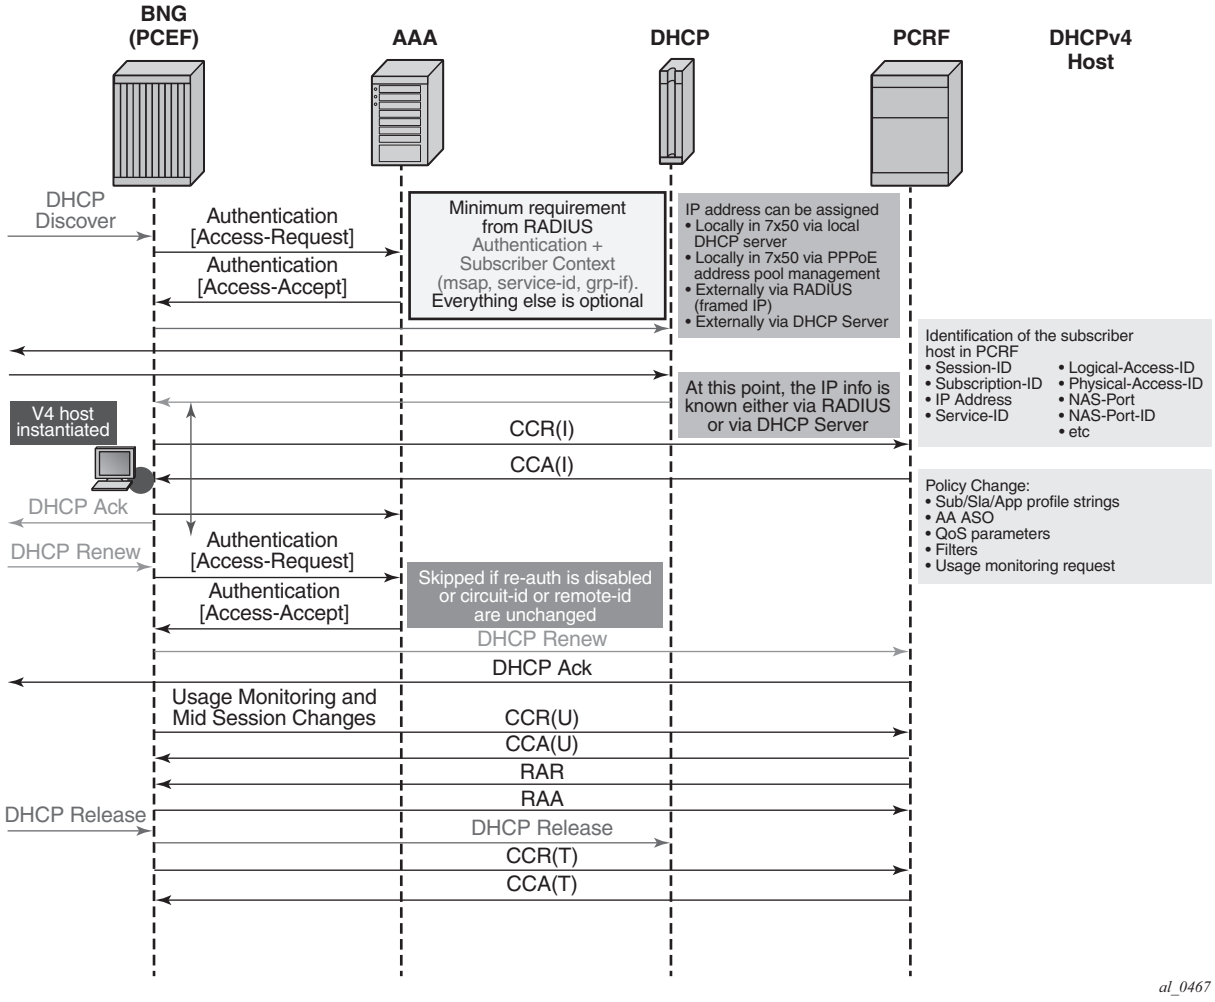

Message flow for PPPoEv4 host is similar. The host is instantiated once the answer from PCRF is received.

However, IPCP negotiation and Gx negotiation (CCR/CCA) are performed in parallel, independently of each other, and therefore the node will not wait for the Gx session to be established before the last IPCP ConfAck is sent (as it is for DHCP ACK).

Once the host is instantiated on the node (after the CCA-i is received or as defined by the fallback action if the PCRF is not available), the Accounting-Start message is sent by the node (assuming that accounting is enabled).

The message flow is shown in Figure 223.

Figure 223 **Message Flow During PPPoEv4 Host Instantiation Phase**

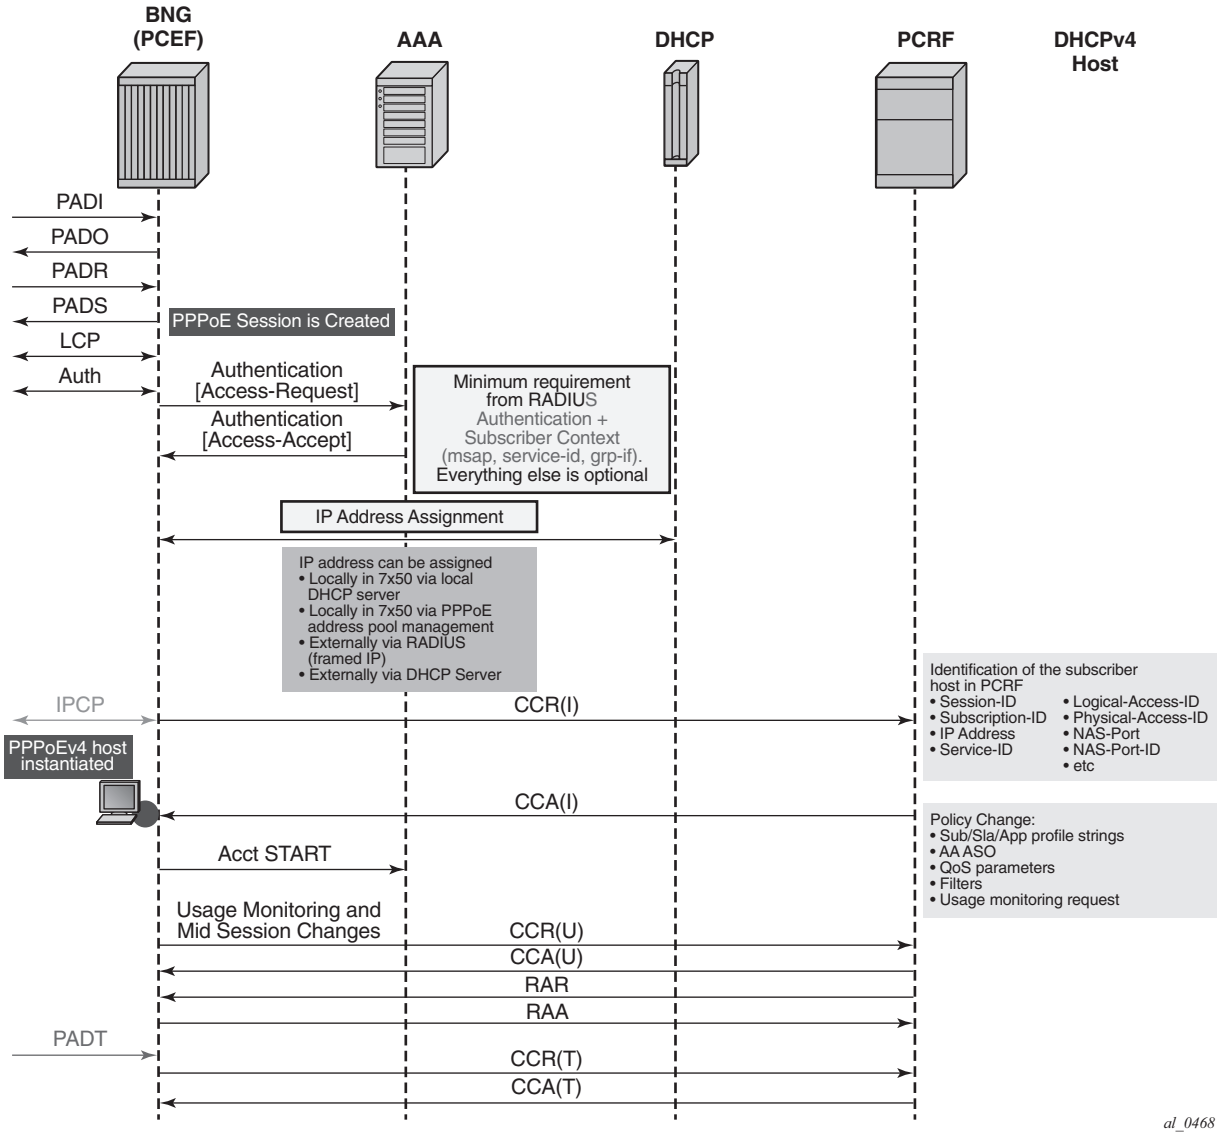

The host is created when the Gx session is established and therefore the subscriber host transitions into the traffic forwarding state once the Gx processing is completed. If the PCRF is unavailable or unresponsive, the host creation or termination is driven by the fallback configuration.

16.5.4.1 Gx and Dual-Stack Hosts

Dual-stack (DS) hosts are treated as a single session from the Gx perspective, regardless whether or not an IPoE session concept is enabled. The difference between the IPoE session and non-IPoE session concept is related to the identification of the hosts for identification and policy management purposes.

If the IPoE session concept is disabled, a dual-stack host is identified by its {MAC,SAP} combination. Consequently, the key used for Gx session creation is based on the same {MAC,SAP} combination.

In certain scenarios, this is not sufficient to differentiate hosts on the same SAP for Gx policy management purposes. As a result. a single Gx session is created for all hosts on that SAP. To avoid this creation, the IPoE session concept must be enabled in SR OS, where a Circuit-id or a Remote-id can be used for further differentiation between the Gx managed entities (subscriber-hosts).

The PCRF submits the policy rules per Gx session. The rules are then be applied to the underlying entity that is managed by this Gx session (single or dual stack IPoE host or IPoE session). For PPPoE deployment, the Gx session is automatically tied to the PPPoE session. A notion of a session is native to PPPoE (there is a session ID in PPPoE), and thus, it is more natural to conceptualize the relationship between a PPPoE session (for single of dual-stack hosts) and a Gx session. By contrast, the concept of a session for IPoE is artificial, and which is why additional consideration is required for IPoE hosts, as described above.

The following example examines a case where IPoE session concept is disabled and consequently the IPoE dual-stack host is tied by a {MAC,SAP} combination.

The CCR-I contains the IP address that was first allocated (meaning, the first IP address that triggered the session creation). The request for the second IP address family triggers (if enabled by configuration) an additional CCR-u that carries the IP address allocation update to the PCRF along with the UE_IP_ADDRESS_ALLOCATE (18) event.

Separately, the CCR-u content mirrors the content of the CCR-i with exception of the pre- allocated IP address or addresses. A single Gx message (CCR-i or CCR-u) carries the update for the DHCPv6 IA-NA+IA-PD and DHCPv6/PPPoE NA+PD address/prefix. This assumes that the NA+PD is requested in a single DHCP message.

Similarly, for the Gx session teardown, the CCR-u messages are sent carrying the UE_IP_ADDRESS_RELEASE event, followed by the CCR-t message.

The message flow is depicted in Figure 224.

Figure 224 **Gx and Dual Stack Session Instantiation**

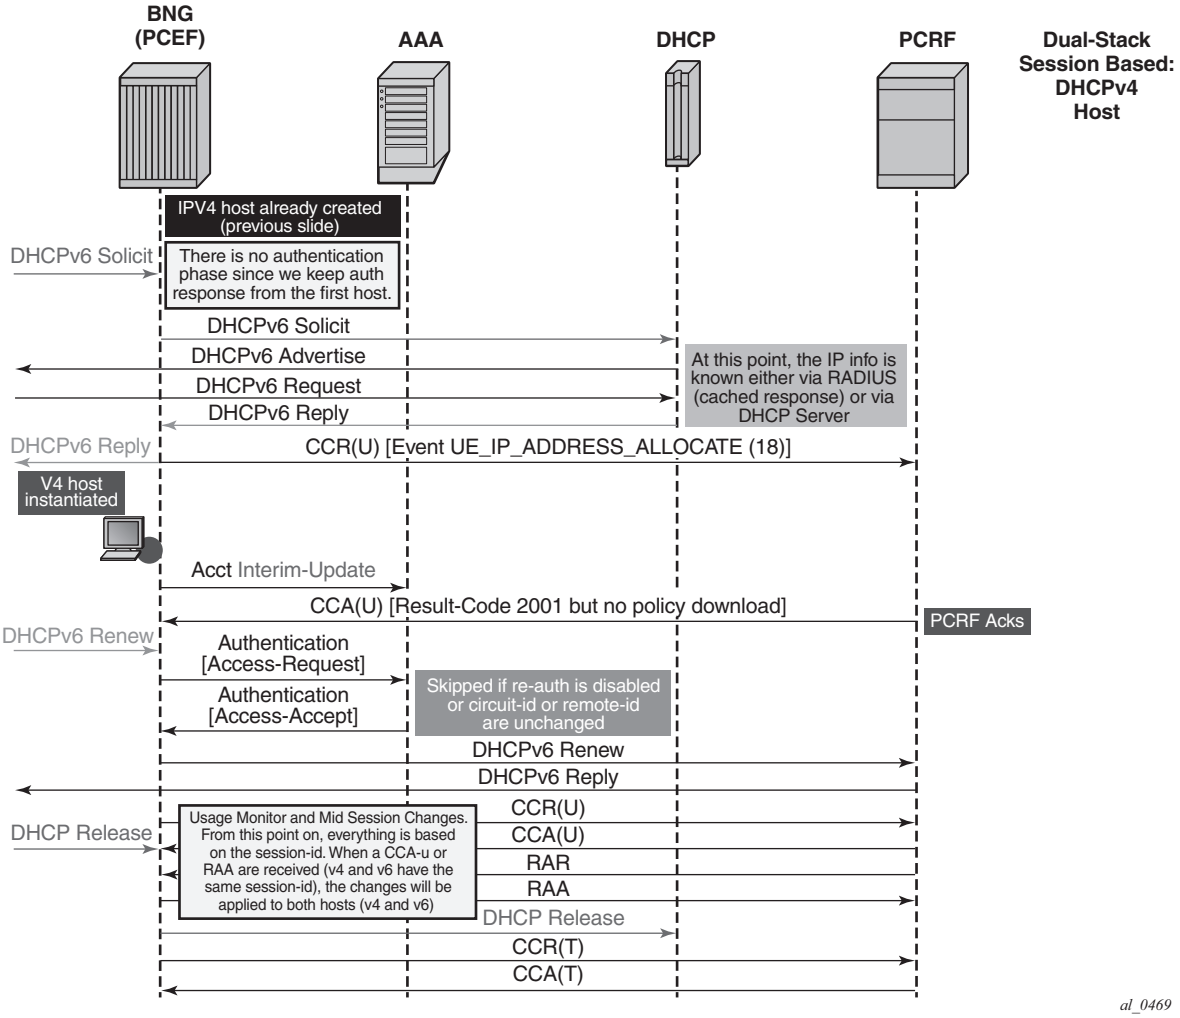

For Dual-Stack PPPoE host, the CCR-i is sent when the first IP address is assigned to the host. In the example in Figure 224, processing of the DHCPv6 Replay and CCR-u messages is performed in parallel. In other words, sending the DHCPv6 Reply message to the client is not delayed until the response from the PCRF is received. The reason being is that the Gx session is already established (triggered by the IPv4 host in our example) and all parameters for IPv4 and IPv6 are already known as received in CCA-i. Then, the CCR-u message is simply a notification message, informing the PCRF about the new IPv6 address/prefix being assigned to an existing client.

16.5.4.2 Gx and PPPoEv6-DHCP

For PPPoE v6 hosts, the IPv6 address is not obtained during the IPCP phase (only interface-id is negotiated). Then, the node waits until the IPv6 address/prefix is allocated to the IPv6 hosts before it sends the CCR-I message. Otherwise, the IP address would not be available in CCR-i. This is shown in Figure 225.

Figure 225 **Gx and PPPoEv6 Host Instantiation**

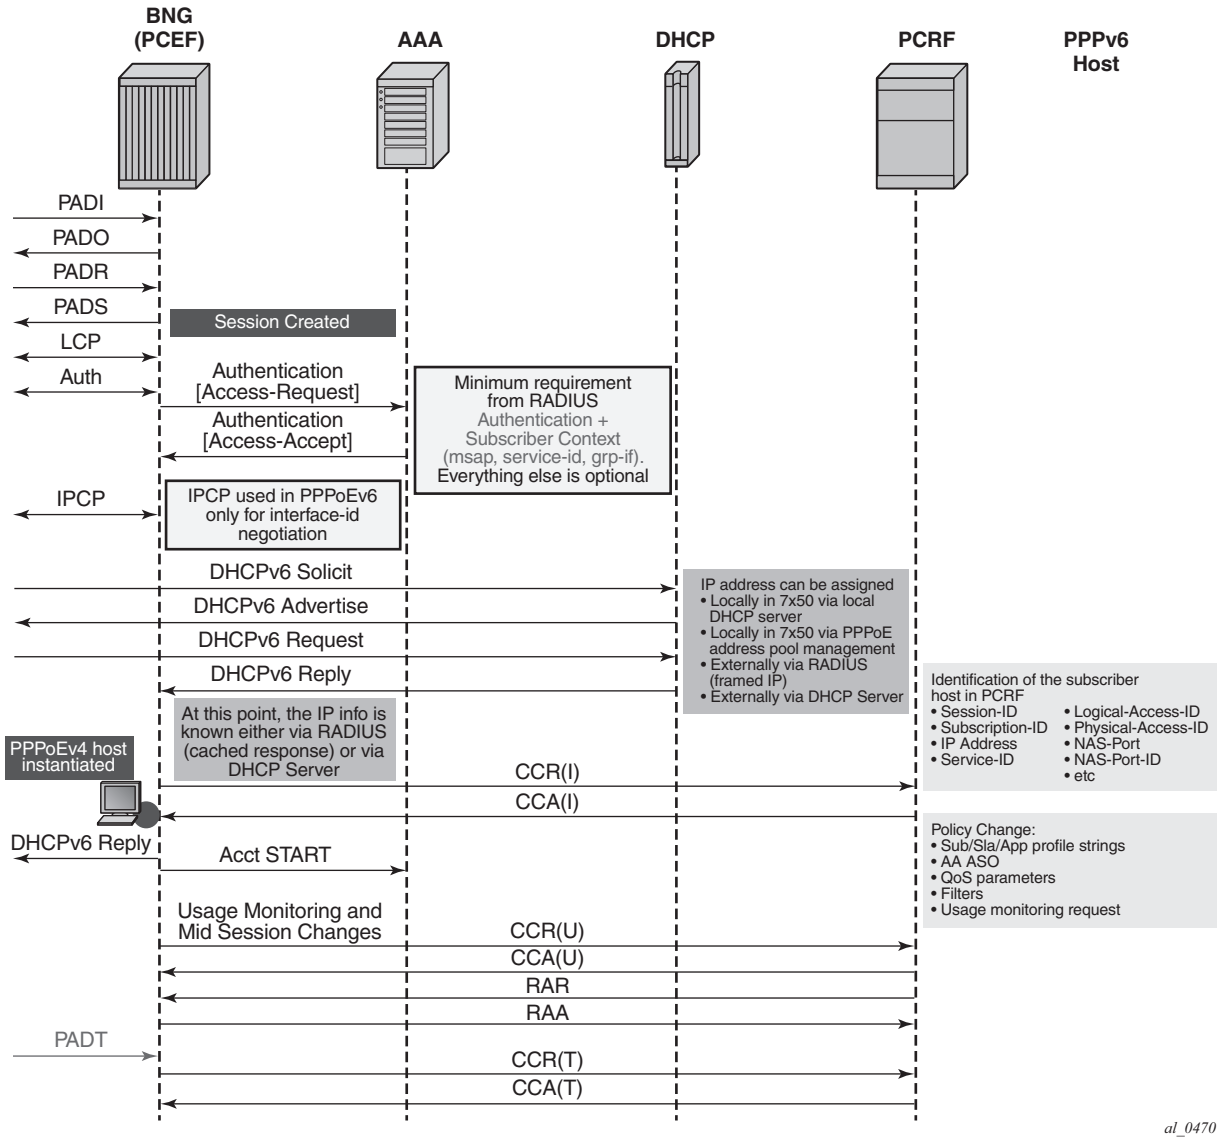

16.5.5 Gx Fallback Function

The Gx fallback functionality refers to the behavior related to the subscriber host instantiation in situations where the PCRF is unresponsive while peering connection(s) are up or the PCRF is unavailable with all peering connections down. This functionality affects only Gx session processing related to CCR-i messages on the node and has no effect on already established Gx sessions.

The fallback behavior can be controlled via local configuration on the node or can be controlled via certain AVPs provided by PCRF.

PCRF-provided AVPs that control fallback behavior are:

- CC-Session-Failover AVP with the following values:
  - FAILOVER_NOT_SUPPORTED
  - FAILOVER_SUPPORTED
- Credit-Control-Failure-Handling AVP with the following values:
  - TERMINATE
  - CONTINUE
  - RETRY_AND_TERMINATE

If the fallback-related AVPs are not provided via PCRF, the node can provide a local configuration option to define the fallback behavior. If the response from the PCRF cannot be obtained, the local configuration can allow the subscriber host to be instantiated with default parameters, or alternatively the local configuration can deny subscriber host instantiation.

PCRF provided AVPs overrule the local configuration.

The local configuration that defines Gx fallback behavior can be found under the following CLI hierarchy:

```
config
    subscr-mgmt
      diam-appl-plcy
        on-failure
          failover {enabled|disabled}
            handling {continue|retry-and-terminate|terminate}
```

The **failover** configuration option (equivalent to CC-Session-Failover AVP) controls whether the secondary peer is used in if the primary peer is unresponsive. The unresponsiveness is determined by the timeout of the previously sent message.

The **handling** configuration option (equivalent to Credit-Control-Failure-Handling AVP) controls whether the subscriber is terminated or instantiated with default parameters if the PCRF is unresponsive.

Table 189 **Handling Options**

	Handling: CONTINUE	**Handling: RETRY-AND-TERMINATE**	**Handling: TERMINATE**
Failover: ENABLED			
Once the message sent to the primary peer times out, the secondary peer (and consecutive peers after that) is attempted. Once the message times out after it has been sent to all available peers, the HANDLING action is examined in order to determine whether to terminate the host instantiation attempt or whether to use default parameters to instantiate the host.	Once the message times out after it has been sent to all available peers, the subscriber host is instantiated with default parameters (if they are configured)	Once the message times out after it has been sent to all available peers, the subscriber host instantiation is terminated.	Once the message sent to the primary peer times out, the subscriber host instantiation is terminated.
Failover: DISABLED			
Once the message sent to the primary peer times out, the HANDLING action is examined in order to determine whether to terminate the host instantiation attempt or whether to use default parameters to instantiate the host.	Once the message sent to the primary server times out, the subscriber host is instantiated with default parameters (if they are configured)	Once the message sent to the primary peer times out, the subscriber host is terminated.	Once the message sent to the primary peer times out, the subscriber host is terminated.

The CCR retransmissions are controlled by the **tx-timer** command under **diameter-application-policy**. See Subscriber Management Command Reference for the description of **retransmission** handling.

If all peers are down (no connections are open), the **handling** action will determine the behavior. If the action is set to **continue**, the subscriber host is immediately instantiated with the default-settings (provided that the defaults are available). In all other action cases, the host instantiation is immediately terminated.

16.5.6 Gx CCR-I Replays

As described in the previous section, the subscriber host can optionally be (configuration controlled) established with default settings (sla-profile, sub-profile, app-profile) if the PCRF is not available to answer the CCR-i. This results in a subscriber-host state mismatch between the node and the PCRF, where the subscriber-host is established on the node but there is no corresponding Gx session established in the PCRF.

In order to resolve this situation, ESM periodically sends CCR-i for the Gx *orphaned subscriber-host* until the response from PCRF is received. The CCR-i is periodically retransmitted every 60 seconds.

16.5.7 Gx CCR-t Replays

Termination of the subscriber-host on the node without termination of the corresponding Gx session in PCRF will result in a state mismatch between the node and the PCRF, whereby the host Gx session is present in the PCRF while it is removed from the node.

Some PCRFs can cope with such out-of-sync condition by periodically auditing all existing Gx sessions. For example, a **probing** RAR can be sent periodically for each active Gx session. The sole purpose of this probing RAR is to solicit a response from the PCEF (node) and provide indication on whether the corresponding Gx session is alive on the node or is vanished. The ‘probing’ RAR can contain an Event-Trigger that is already applied on the node, or if none is applied, then the Event-Trigger can contain NO_EVENT_TRIGGER. In either case the ‘probing’ RAR will not cause any specific action to be taken in the node and it is used only to solicit reply from PCRF.

To minimize the impact on performance, **probing** RARs are sent infrequently; therefore, it may take days to discover a stale Gx session on the PCRF. The node supports a mechanism that can clear the stale session in the PCRF sooner. It does this by replaying CCR-t messages until the proper response from the PCRF is received (CCA-t). The CCR-t messages are replayed up to 24 hours. Both the interval at which the CCR-t messages are replayed and the max-lifetime period are configurable. If the max-lifetime period expires before a valid answer is received, the CCR-t is deleted and a log is generated. The log contains Gx session ID.

The T-bit (retransmission bit) is set in all replayed CCR-t messages.

The following command clears all orphaned sessions on the node for a specified Diameter application policy:

```
clear subscriber-mgmt diameter-session ccrt-replay diameter-application-policy gx-
policy-name sessions
```

16.5.7.1 RAR and CCR-t Replay

Certain scenarios will make possible that the PCRF sends a RAR message to an orphaned Gx session running CCR-t replays on the node. The ESM host associated with the orphaned Gx session does not exist and therefore RAR cannot be applied.

In this scenario, the node will reply with RAA carrying Result-Code= DIAMETER_UNKNOWN_ SESSION_ID (5002).

16.5.7.2 CCR-t Replay And Multi-Chassis Redundancy

The first CCR-t reply for each Gx session is synchronized, but the consecutive CCR-t replays for the same Gx sessions are not synchronized. Once the answer (CCA-t) is received, the CCR-t replay is terminated and this event (deletion of CCR-t replay) is synchronized to the other node.

CCR-t replays are sent from the node that was in SRRP active state at the time when the CCR-t was initiated. They continue to be sent from the same node even if the SRRP is switched over in the meantime.

This entire process can be thought of as if the CCR-t initiating node (active SRRP) armed its MCS peer with CCR-t replay for a given Gx session. This occurs at the very beginning, when a CCR-t replay is first initiated for a given Gx session. The armed node will stay silent until the MCS peer that is actively sending CCR-t replays for a given Gx session, reboots. Only when the MCS peer reboots, the armed node will start sending CCR-t replays for a given Gx session in the following fashion: first message is sent with cleared T-bit, followed by replays at the configured replay interval and a fresh 24 hour lifetime.

16.5.7.3 CCR-t Replay And High Availability

On CPM switchover, the newly active CPM resumes the CCR-t replay with T-bit set until the lifetime, which is synchronized between the CPMs, expires.

16.5.8 Automatic Updates for IP Address Allocation/De-allocation

During the subscriber-host setup phase, the first allocated IP address is sent in the CCR-i message from the node to the PCRF.

Each subsequent IP address allocation or de-allocation for the same host can optionally trigger a CCR-u, notifying the PCRF of the IP address allocation/de-allocation event.

This behavior can be enabled via the following CLI command:

```
configure
    subscriber-mgmt
      diameter-application-policy <pol-name>
        gx
          [no] report-ip-addr-event
```

The IP address allocation/de-allocation event driven CCR-u message will carry the respective event code [UE_IP_ADDRESS_ALLOCATE(18) or UE_IP_ADDRESS_RELEASE(19)] along with the corresponding IP address.

The IP address allocation/de-allocation events are applicable to the following addresses:

- Framed-IP-Address (AVP Code 8) IPv4
- Framed-IPv6-Prefix (AVP Code 97) SLAAC

- Delegated-IPv6Prefix (AVP Code 123) IA-PD
- Alc-IPv6-Address (AVP Code 1023) IA-NA

These event-codes will only be sent in CCR-u messages and not in CCR-i and CCR-t messages (when the host is instantiated and terminated).

Examples:

- IPv6 attachment request arrives with two IP addresses: IA-NA and IA-PD. This is a new host. CCR-i is generated with two IP addresses included (IA-NA and IA-PD, assuming that request for their allocation is carried in the same DHCPv6 message).
- Afterward, the attachment request for an IPv4 address arrives on the same host. CCR-u is generated with the event UI_IP_ADDRESS_ALLOCATE and corresponding AVP (framed-address) is sent to the PCRF. No IP address other than this new IPv4 address is sent.
- RAR is received for the (any) policy change. The node will reply with RAA and it will contain all three IP addresses (AVPs) that have been allocated to the host.

If the IP address notification event is enabled, the node-originated Gx message will carry all known IP addresses/prefixes associated with the subscriber-host (Gx session), unless those messages contain one of the two event codes:

UE_IP_ADDRESS_ALLOCATE(18) or UE_IP_ADDRESS_RELEASE(19).

If one of those two events is present in the Gx message, the IP address/prefix carried in that message is only relevant to the event contained in the message (address/prefix allocated or released).

If the IP address notification event is disabled, the node will only send the IP address from the first host. This IP address is included in all messages related to the Gx session. If this IP address is removed (de-allocated) mid-session from the dual-stack host, the node will stop advertising it, or any other address, from Gx messages for that particular session.

16.5.9 DHCPv4/v6 Re-Authentication and RADIUS CoA Interactions With Gx

If re-authentication for DHCPv4/v6 hosts is enabled, any policy changes that may be submitted during re-authentication (for example sla-profile update via Access-Accept) will overwrite the one previously applied, regardless of the source of the policy update. For example, if that the Gx policy is applied to a subscriber host via RAR (mid-session policy update) and then some time later an overlapping policy with

different values is submitted via RADIUS or LUDB during the re-authentication phase, the RADIUS/LUDB submitted policy will overwrite the one applied via Gx. In other words, the origin of the current policy in effect is not maintained internally in the system and therefore the overlapping policy update cannot be prioritized according to the source of the policy.

The following guidelines should be followed if the policy is provided via Gx:

- If LUDB access is enabled, there should be no overlap between the LUDB provided parameters and Gx provided parameters. LUDB is accessed during every DHCP lease renew process and consequently parameters configured via LUDB would overwrite parameters provided by Gx.
- If LUDB access is enabled, there should be no overlap between the LUDB provided parameters and Gx provided parameters. LUDB is accessed during every DHCP lease renew process and consequently parameters configured via LUDB would overwrite parameters provided by Gx.
- If re-authentication is enabled, there should be no overlap between the RADIUS provided parameters and Gx provided parameters. With re-authentication enabled, RADIUS is contacted during every DHCP lease renew process and consequently parameters configured via RADIUS would overwrite parameters provided by Gx.

These guidelines are not applicable for PPPoE subscriber-hosts since re-authentication cannot be enabled for PPPoE hosts. Consequently, LUDB or RADIUS parameters cannot override Gx provided parameters.

Coexistence of RADIUS CoA and Gx for the same host is allowed. The two policy change mechanisms are independent of each other and as such they can override each other. For example, if the RADIUS CoA for policy change for the host is received, the policy is updated but the PCRF (Gx) will not be notified of the change. If both policy management mechanisms are deployed simultaneously, then it is the operator’s responsibility to synchronize the actions between the two.

16.5.10 Gx, ESM and AA

Although the ESM subscriber and the AA subscriber are two separate instantiations on the node, their policy management and Usage-Monitoring are handled uniformly through a single Gx session.

16.5.10.1 ESM Subscriber-Host vs AA Subscriber

Since ESM and AA modules are part of integrated service offering (ESM with residential AA on the same node), they share the same subscriber-id string. However, Gx interface in ESM is primarily applicable to hosts (basic entity to which policy is applied) while AA has no awareness of hosts. AA is only aware of subscribers (which is, in broader terms, a collection of hosts within a residence). Refer to the *7450 ESS, 7750 SR, and VSR Multiservice Integrated Service Adapter Guide* for details on Application Assurance concepts.

16.5.10.2 AA Subscriber State

AA subscriber state must exist for App-profiles and ASO overrides to be applied.

The app-profile for the aa-sub is applied explicitly by a CCR-i or RAR message with an AVP AA-App-Profile-Name.

App-profiles interact with ASO characteristics in this way:

- The AA-App-Service-Options AVP within the app-profile assignment is optional at subscriber instantiation time and may be used later to modify the policy.
- The newly submitted AA-App-Profile-Name AVP will overwrite the one that is already applied. Any ASO AVPs that is received within the Gx message is applied.

Note: If an app-prof AVP is present, even if it is the same app-profile as currently applied, all previous ASO override policies are removed for the sub.

The state of the subscriber policy attributes is modified by ASO AVPs in this way:

- The app profile can define one or more ASO characteristics attributed to a subscriber
- If there are multiple ASO AVPs for the same characteristic in the message, the first one will take effect.
- There is no explicit delete of ASO overrides (the PCRF can always resend or change the app-profile in order to delete all overrides).

16.5.11 Policy Management via Gx

A policy change can be implicitly requested by the node at IP-Can session establishment time via the CCR-i command. The node will supply user identification attributes to the PCRF so that the PCRF can identify rules to be applied. However, the node will not explicitly request specific policy update, for example via Event-Trigger = RESOURCE_MODIFICATION_REQUEST.

Another way to request a policy update on the node is via the **rar** command in a PUSH model.

Gx policies on the node can be enforced via these three mechanisms:

- Gx-based overrides — This refers to subscriber related overrides (sub/sla/aa-profile, subscriber-id, QoS, filter, category-map, and so on).
- PCC rules or IP-criterion based rules which are fully constructed Policy and Charging Control (PCC) rules with multiple qos/filter actions per rule and its own traffic classification.
- NAS filter entry inserts via Gx — Basic permits or denies entries, similar to ACL filter entries.

16.5.12 Gx-Based Overrides

Gx-based overrides refer to the activation or modification of the existing subscriber host- related objects on the node.

Subscriber host-related objects are shown in Figure 226. A subscriber represents a residence or home and it is identified by Subscriber-Id string on the node. A subscriber on the node can consist of multiple hosts in a bridged home environment or a single host in a routed home environment.

Figure 226 **ESM Objects Managed via Various Policies and Profiles**

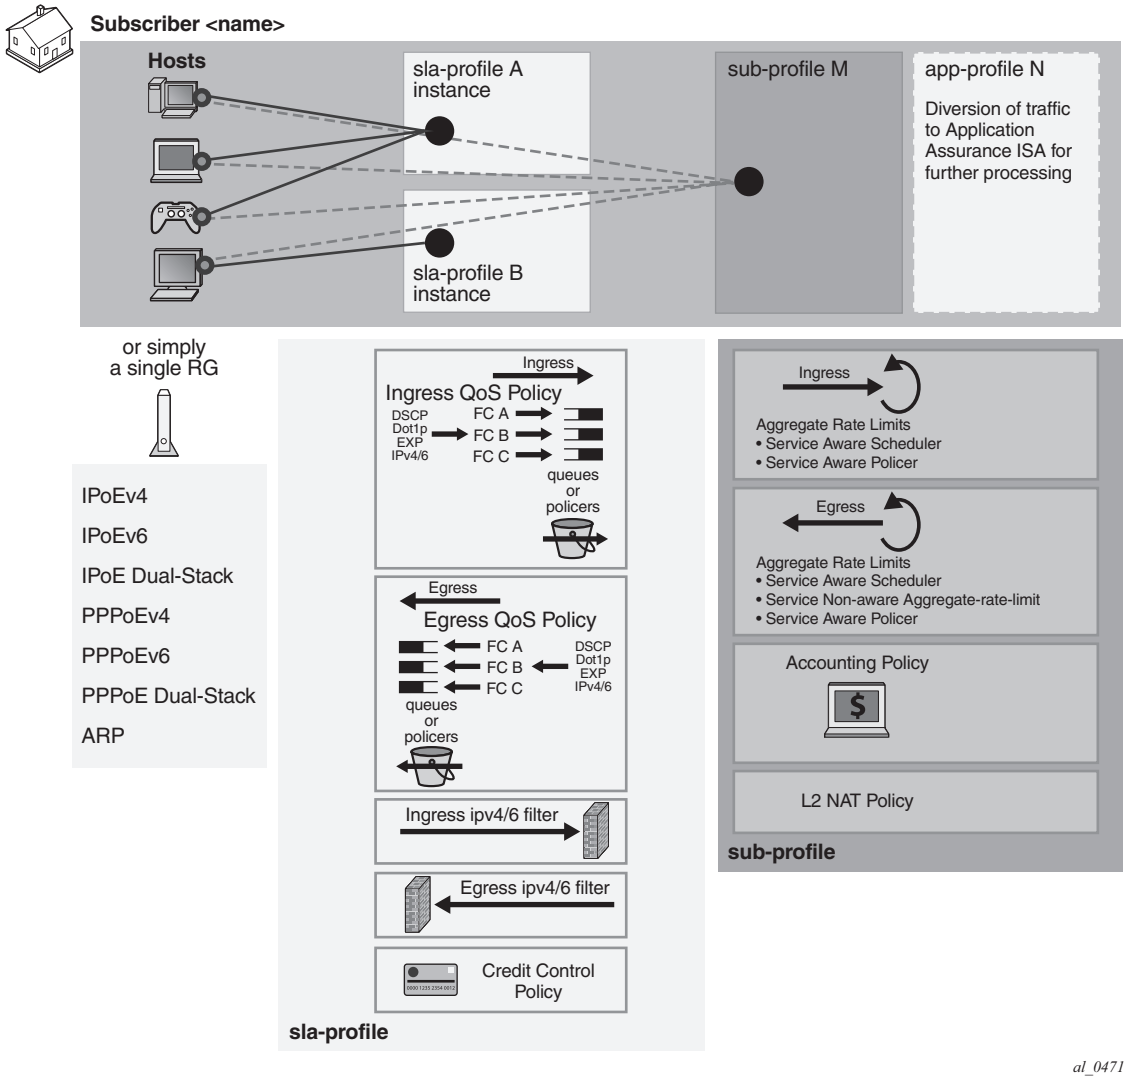

The two basic concepts in ESM context are sla-profile with its associated objects and sub-profile with its associated objects.

- Sla-profile defines a service level (rates, queues, filters) for a group of hosts sharing the same sla-profile instance within the subscriber. There can be multiple sla-profile instances per subscriber.
- Sub-profile defines aggregate rate of the subscriber along with accounting policy. There is only one sub-profile per subscriber.

16.5.12.1 Instantiation of Gx Overrides

For a list of Gx related AVPs supported on the node, refer to the *7750 SR Gx AVPs Reference Guide*.

Gx overrides are installed via Charging-Rule-Install AVP (for ESM or AA) or ADC-Rule-Install AVP (for AA only – 3GPP Release 11) sent from the PCRF towards the node.

AVP Format:

```
Charging-Rule-Install ::=       < AVP Header: 1001 >
                                       *[ Charging-Rule-Definition ]
                                       *[ Charging-Rule-Name ]
                                       *[ AVP ]

ADC-Rule-Install ::=       < AVP Header: 1092 >
                                       *[ ADC-Rule-Definition ]
                                       *[ ADC-Rule-Name ]
                                       *[ AVP ]
```

Every Gx override must have a Charging-Rule-Name (ESM) or ADC-Rule-Name (AA - 3GPP Release 11 and Release 12) associated with it. This is important in order to return the override status from the node to the PCRF upon the override instantiation.

The objects (subscriber-hosts) to which the new overrides are applied must exist on the node; otherwise, the override installation will fail.

The parameters defining a new override will simply replace the existing parameters that are already applied to the subscriber-host, without the need to remove the previously installed parameters.

There are four types of overrides that are currently supported via Gx:

- ESM string overrides (sla/sub/app-profiles, subscriber-id)
- Update of subscriber host QoS information (queue rate change)
- Filter override for the subscriber host (including one-time http redirect)
- Category-map installation/override
- HTTP redirect override within the filter.

A Charging-Rule-Name AVP within the Charging-Rule-Install grouped AVP can have several meanings.

- It can directly reference a preconfigured filter in the system.
- It can directly reference a preconfigured subscriber profile (sla and sub).
- It can directly reference a preconfigured category map.

- It can directly reference an ESM string (such as an inter-destination-string used to associate the subscriber host with a vPort construct).
- Or, if there is an HTTP redirect override, it serves as a special container within which the HTTP redirect URL is defined. Note that HTTP redirect override is not to be confused with PCC rule-based HTTP redirect.

In all of the above cases, the existing objects applied to the subscriber-host is replaced with the referenced object.

It is important to distinguish two locations for invoking the Charging-Rule-Name AVP for overrides:

1. Directly under the Charging-Rule-Install AVP – Then, the Charging-Rule-Name will reference the predefined structures (profiles, filter-ids, cat-maps, and so on) on the node. The type of the structure is contained within the Charging-Rule-Name AVP in the form of a reserved keyword that has to be prepended to the identifier of structure:

   Filter installation/overrides:

   – Charging-Rule-Name = Ingr-v4:<id>
   – Charging-Rule-Name = Ingr-v6:<id>
   – Charging-Rule-Name = Egr-v4:<id>
   – Charging-Rule-Name = Egr-v6:<id>
   – Charging-Rule-Name = In-Othr-v4:<id> (othr - one-time-http-redirect)
   – Charging-Rule-Name = In-Othr-v6:<id> (othr - one-time-http-redirect)

   Subscriber-Id override:

   – Charging-Rule-Name = Sub-Id:sub-id-string

   Profile installation/overrides:

   – Charging-Rule-Name = Sla-Profile:sla-profile-name
   – Charging-Rule-Name = Sub-Profile:sub-profile-name

   Inter-destination string override:

   – Charging-Rule-Name = Inter-Dest:inter-dest-string

   Usage-Monitoring:

   – Charging-Rule-Name = Cat-Map:category-map-name

   AA:

   – Charging-Rule-Name = AA-UM:<string-name>
   – Charging-Rule-Name=AA-Functions:<string-name>

In summary, the reserved prefixes ingr-v4:,ingr-v6:, egr-v4:, egr-v6:, in-othr-v4:, in-othr-v6:, sub-id:, sla-profile:, sub-profile:, inter-dest:, cat-map:, aa-um: and aa-functions: have special meaning within the Charging-Rule-Name AVP on the node.

One exception to this is HTTP redirect override, which is described separately.

2. Under the Charging-Rule-Install — Charging-Rule-Definition AVP. In this case, the override itself is not pre-provisioned on the node but instead, directly defined in the Charging-Rule-Definition. Part of the override definition is the name assignment via Charging-Rule-Name AVP. The Charging-Rule-Name AVP is used to report on the override status.

   For example, the Charging-Rule-Name AVP for QoS overrides is an arbitrary name. This AVP is part of Charging-Rule-Definition AVP in which QoS-Information is provided. Such Charging-Rule-Name is used to report errors related to instantiation of the override.

   Another example of the override defined within the Charging-Rule-Definition AVP is HTTP redirect override:

```
Charging-Rule-Name = v4-http-url:<name>
Redirect-Information ::= < AVP Header: 1085 >

[ Redirect-Address-Type ]
[ Redirect-Server-Address ]
```

and

```
Charging-Rule-Name = v6-http-url:<name>
Redirect-Information ::= < AVP Header: 1085 >
[ Redirect-Address-Type ]
 [ Redirect-Server-Address ]
```

The ADC-Rule-Name AVP within the ADC-Rule-Install grouped AVP handles application policy related processing (AA). This AVP is applicable under the ADC-Rule-Install—ADC-Rule-Definition AVP. In this case the ADC rule itself is not pre-provisioned on the node but instead directly defined in the ADC-Rule-Definition. In AA, such rule definition can define AA overrides that are applied to the subscriber. In other words, the existing objects for the subscriber are replaced with the ones referenced in the rule. Part of the ADC rule definition is the ADC rule name assignment via ADC-Rule-Name AVP. The ADC-Rule-Name defined in such manner is used to report on the rule status.

The AA-Functions: prefix in the ADC rule name is reserved for ADC rule definitions applicable to AA-functions (namely, app-profile and ASOs):

ADC-Rule-Name = AA-Functions:aa-rule-name

Then the aa-rule-name is an arbitrary name that is used in rule status reporting.

If that ADC-Rule-Name is used in AA Usage-Monitoring, the “AA-Functions:” prefix must not be present (Usage-Monitoring in AA is covered in details in the 7450 ESS, 7750 SR, and VSR Multiservice Integrated Service Adapter Guide).

Note: AA-Function AVP and AA-Usage-Monitoring cannot co-exist in the same ADC rule.

The Charging-Rule-Definition AVP (AVP code 1003, 3GPP 29.212 §5.3.4) is of type Grouped, and it defines the override sent by the PCRF to the node.

The Charging-Rule-Name in this AVP can be arbitrarily set and it is used to uniquely identify the override in error reporting.

```
Charging-Rule-Definition ::= < AVP Header: 1003 >
                                { Charging-Rule-Name }
                                [ QoS-Information ]
                                [ Nas-Filter-Rule]
                                [ Alc-NAS-Filter-Rule-Shared]
                                *[ AVP ]
```

The ADC-Rule-Definition AVP (AVP code 1094, 3GPP 29.212 §5.3.87) is of type Grouped, and it defines the ADC override sent by the PCRF to the node. The ADC-Rule-Name AVP within the ADC-Rule-Definition AVP uniquely identifies the ADC policy rule and it is used to reference to a policy rule in communication between the node and the PCRF within one IP CAN session.

```
ADC-Rule-Definition ::= < AVP Header: 1094 >
             { ADC-Rule-Name }
             [AA-Functions]
             *[ AVP ]
```

16.5.12.2 HTTP Redirect Override

HTTP redirect override submitted via the Gx interface will override the current URL string defined in the filter that is currently applied to the subscriber-host.

The override is implemented through the standard Redirect-Information AVP nested within the Charging-Rule-Definition (CRD) AVP.

IPv4

```
Charging-Rule-Install
   Charging-Rule-Definition
      Charging-Rule-Name = v4-http-url:<name>
         Redirect-Information ::= < AVP Header: 1085 >
           [ [ Redirect-Address-Type ]
```

```
          [ [ Redirect-Server-Address ]
```

IPv6

```
Charging-Rule-Install
   Charging-Rule-Definition
      Charging-Rule-Name = v6-http-url:<name>
         Redirect-Information ::= < AVP Header: 1085 >
           [ Redirect-Address-Type ]
           [ Redirect-Server-Address ]
```

The keywords **v4-http-url** and **v6-http-url** are special keywords that must be part of the CRN AVP. These keywords can be followed by an arbitrary string *name*.

The purpose of the <name> string in the Charging-Rule-Name (CRN) AVP is for the PCRF to differentiate between different HTTP redirect overrides. However, the *name* string in the context of the **http url host override** command in a filter has no meaning on the node, and therefore it is ignored. This means that there can be only one HTTP redirect override per host and per address family on a node.

The outcome of this Gx directive (Redirect-Information AVP without the Flow-Information AVP within the Charging-Rule-Definition AVP) is the override of the HTTP redirect URL in the currently applied subscriber-host filter. The filter definition must explicitly allow overrides via the **allow-radius-override** keyword.

As long as the override rule is present in the system (meaning, it has been submitted via the Gx and has not been explicitly removed since), the override will try to enforce itself when both of the following two conditions are met:

- A filter with the action **http-redirect** is applied to the subscriber-host. This is valid whether the filter was there when the rule was submitted, or whether the filter was applied after the rule was submitted.
- The filter definition allows **http-redirect** overrides (filter with the **allow-radius-override** keyword).

If the above conditions are not met, the override is accepted (the node responds with RAA=OK) and stored by the system, although it will not be applied until the above conditions are met.

For the HTTP URL host, the CRD directive must not contain any flow information or any other action besides the Redirect-Information AVP. Otherwise, the Diameter encoding will fail and an error response is generated for RAR while CCR-I is silently dropped.

16.5.12.3 Removal of Overrides

With the exception of HTTP redirect override, overrides cannot be removed by the Charging-Rule-Remove AVP. They can only be overridden, and consequently the Charging-Rule-Remove AVP is ignored. It is ignored only for regular overrides and not for PCC rules (see PCC Rules) or for HTTP redirect override.

An HTTP redirect override can be removed whether it is active (a filter with HTTP redirect action is applied) or inactive (a filter without HTTP redirection is applied).

```
Charging-Rule-Remove ::= < AVP Header: 1002 >
   Charging-Rule-Name = v4-http-url:<name>
   Charging-Rule-Name = v6-http-url:<name>
```

The *name* string in the CRN AVP is ignored in the context of HTTP redirect override. This means that the removal of HTTP redirect override with any name will remove the currently installed HTTP redirect override.

Similarly, the installation of the HTTP redirect override will replace any currently installed HTTP redirect override, regardless of the *name* string (implicit removal of the current HTTP redirect override, followed by the installation of the new one).

The node will reply with RAA=OK if a properly formatted Charging-Rule-Remove directive with any name is received for HTTP redirect override.

16.5.12.4 Examples of Gx Overrides

The instantiation of HTTP redirect overrides via the Gx can be summarized as:

- The Central AVP for the Gx overrides in SR OS is the Charging-Rule-Install AVP. Multiple overrides can be submitted to an node via a single Charging-Rule-Install AVP, or each Gx override can be submitted via its own Charging-Rule-Install AVP.
- Gx override is identified by the Charging-Rule-Name AVP. This AVP is also used to report on the status of Gx override. The Charging-Rule-Name can reference a pre-configured construct on the node (profiles, cat-maps, filters) or it can be assigned by PCRF to identify the PCRF defined override (QoS policy modifications, HTTP redirect AA ASO modifications, and so on).
- With the exception of the HTTP redirect override, overrides cannot be removed by the Charging-Rule-Remove AVP. They can only be overridden. The Charging-Rule-Remove AVP is ignored for Gx overrides.

The following is an example of Gx override instantiation where all Gx overrides are submitted under a single Charging-Rule-Install AVP. The AVPs in this example can be included in the CCA-i, CCA-u or RAR messages sent from the PCRF.

The outcome of the override is the following:

- The subscriber-id is overridden with the new name ‘residence-1’.
- New ingress v4 and egress v6 filters are installed.
- New sub-profiles and sla-profiles are applied to the subscriber host.
- The subscriber host is associated with the vPort named *vport-AN-1*.
- The existing HTTP redirect URL is overwritten with a new URL.
- Characteristics for queue 5 and 7 are overridden.
- The egress rate limit for the subscriber is overridden.
- ASOs are overridden.

```
Charging-Rule-Install ::= <AVP Header: 1001>
Charging-Rule-Name <AVP Header: 1005> = "sub-id:residence-1"
Charging-Rule-Name <AVP Header: 1005> = "ingr-v4:7"
Charging-Rule-Name <AVP Header: 1005> = "eggr-v6:5"
Charging-Rule-Name <AVP Header: 1005> = "Sub-Profile:prem"
Charging-Rule-Name <AVP Header: 1005> = "Sla-Profile:voip+data"
Charging-Rule-Name <AVP Header: 1005> = "Inter-Dest:vport-AN-1"

Charging-Rule-Definition <AVP Header: 1003>
Charging-Rule-Name <AVP Header: 1005> "v4-http-url:http-redir-1"
Redirect-Information < AVP Header: 1085 >
  Redirect-Address-Type < AVP Header: 433 > = 2 -> URL type
    Redirect-Server-Address < AVP Header: 435 > =
"http://www.operator.com/portal.php?mac=$MAC"


Charging-Rule-Definition <AVP Header: 1003>
Charging-Rule-Name <AVP Header: 1005> = "premium-service"
QoS-Information <AVP Header: 1016>
Alc-Queue <AVP Header; vnd ALU; 1016>
 Alc-Queue-id <AVP Header; vnd ALU; 1007> = 5
 Max-Requested-Bandwidth-UL <AVP Header: 516> = 10000
         Max-Requested-Bandwidth-DL <AVP Header: 515> = 100000
         Guaranteed-Bitrate-UL <AVP Header: 1026> = 5000
             Guaranteed-Bitrate-DL <AVP Header: 1027> = 50000
 Alc-Committed-Burst-Size-UL <AVP Header; vnd ALU; 1008> = 1000
     Alc-Maximum-Burst-Size-UL <AVP Header; vnd ALU; 1009> = 2000
     Alc-Committed-Burst-Size-DL <AVP Header; vnd ALU; 1010> = 1000
     Alc-Maximum-Burst-Size-DL <AVP Header; vnd ALU; 1011> = 2000

 Alc-Queue <AVP Header; vnd ALU; 1006>
Alc-Queue-id <AVP Header; vnd ALU; 1007> = 7
 Max-Requested-Bandwidth-UL <AVP Header: 516> = 10000
         Max-Requested-Bandwidth-DL <AVP Header: 515> = 100000
         Guaranteed-Bitrate-UL <AVP Header: 1026> = 5000
```

```
        Guaranteed-Bitrate-DL <AVP Header: 1027> = 50000
 Alc-Committed-Burst-Size-UL <AVP Header; vnd ALU; 1008> = 1000
     Alc-Maximum-Burst-Size-UL <AVP Header; vnd ALU; 1009> = 2000
     Alc-Committed-Burst-Size-DL <AVP Header; vnd ALU; 1010> = 1000
     Alc-Maximum-Burst-Size-DL <AVP Header; vnd ALU; 1011> = 2000

Alc-Sub-Egress-Rate-Limit <AVP Header; vnd ALU; 1016> = 10000



ADC-Rule-Install ::= <AVP Header: 1092>
ADC-Rule-Definition <AVP Header: 1094>
ADC-Rule-Name <AVP Header: 1096> = "AA-Functions:apps"
AA-Functions
AA-App-Profile-Name = "apps-prof"
AA-App-Service-Options
AA-App-Serv-Options-Name = "bitttorent"
AA-App-Serv-Options-Value = "low-prio-1mbps"
AA-App-Service-Options
AA-App-Service-Options-Name =  "ftp"
AA-App-Service-Options-Value = "hi-prio"
```

In the following example, all Gx overrides are submitted via a separate Charging-Rule-Install AVP:

```
Charging-Rule-Install ::= <AVP Header: 1001>
     Charging-Rule-Name <AVP Header: 1005> = "sub-id:residence-1"

Charging-Rule-Install ::= <AVP Header: 1001>
     Charging-Rule-Name <AVP Header: 1005> = "ingr-v4:7"

Charging-Rule-Install ::= <AVP Header: 1001>
     Charging-Rule-Name <AVP Header: 1005> = "eggr-v6:5"

Charging-Rule-Install ::= <AVP Header: 1001>
     Charging-Rule-Name <AVP Header: 1005> = "Sub-Profile:prem"

Charging-Rule-Install ::= <AVP Header: 1001>
     Charging-Rule-Name <AVP Header: 1005> = "Sla-Profile:voip+data"

Charging-Rule-Install ::= <AVP Header: 1001>
     Charging-Rule-Name <AVP Header: 1005> = "Inter-Dest:Gx-inserted-string"

Charging-Rule-Install ::= <AVP Header: 1001>
   Charging-Rule-Definition <AVP Header: 1003>
      Charging-Rule-Name <AVP Header: 1005> "v4-http-url:http-redir-1"
         Redirect-Information < AVP Header: 1085 >
            Redirect-Address-Type < AVP Header: 433 > = 2  ? URL type
            Redirect-Server-Address < AVP Header: 435 > =
                  "http://www.operator.com/portal.php?mac=$MAC"
Charging-Rule-Install ::= <AVP Header: 1001>
     Charging-Rule-Definition <AVP Header: 1003>
          Charging-Rule-Name <AVP Header: 1005> = "premium-service"
          QoS-Information <AVP Header: 1016>
               Alc-Queue <AVP Header; vnd ALU; 1016>
                 Alc-Queue-id <AVP Header; vnd ALU; 1007> = 5
                         Max-Requested-Bandwidth-UL <AVP Header: 516> = 10000
```

```
                    Max-Requested-Bandwidth-DL <AVP Header: 515> = 100000
            Guaranteed-Bitrate-UL <AVP Header: 1026> = 5000
            Guaranteed-Bitrate-DL <AVP Header: 1027> = 50000
                        Alc-Committed-Burst-Size-
UL <AVP Header; vnd ALU; 1008> = 1000
                    Alc-Maximum-Burst-Size-UL <AVP Header; vnd ALU; 1009> = 2000
                    Alc-Committed-Burst-Size-DL <AVP Header; vnd ALU; 1010> = 1000
                Alc-Maximum-Burst-Size-DL <AVP Header; vnd ALU; 1011> = 2000

                 Alc-Queue <AVP Header; vnd ALU; 1006>
                        Alc-Queue-id <AVP Header; vnd ALU; 1007> = 7
                        Max-Requested-Bandwidth-UL <AVP Header: 516> = 10000
            Max-Requested-Bandwidth-DL <AVP Header: 515> = 100000
            Guaranteed-Bitrate-UL <AVP Header: 1026> = 5000
            Guaranteed-Bitrate-DL <AVP Header: 1027> = 50000
                        Alc-Committed-Burst-Size-
UL <AVP Header; vnd ALU; 1008> = 1000
                Alc-Maximum-Burst-Size-UL <AVP Header; vnd ALU; 1009> = 2000
                Alc-Committed-Burst-Size-DL <AVP Header; vnd ALU; 1010> = 1000
                Alc-Maximum-Burst-Size-DL <AVP Header; vnd ALU; 1011> = 2000

                 Alc-Sub-Egress-Rate-Limit <AVP Header; vnd ALU; 1016> = 10000

ADC-Rule-Install ::= <AVP Header: 1092>
     ADC-Rule-Definition <AVP Header: 1094>
          ADC-Rule-Name <AVP Header: 1096> = "AA-Functions:apps"
          AA-Functions
               AA-App-Profile-Name = "apps-prof"
               AA-App-Service-Options
                    AA-App-Service-Options-Name = "bitttorent"
                    AA-App-Service-Options-Value = "low-prio-1mbps"
               AA-App-Service-Options
                    AA-App-Service-Options-Name =  "ftp"
                    AA-App-Service-Options-Value = "hi-prio"
```

Gx overrides (QoS rates, sub/sla-profiles, filters, and so on) can be examined individually with subscriber specific operational commands. In the example below, fields in bold can be overridden.

```
show service active-subscribers detail
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber residence-1 (prem)
-------------------------------------------------------------------------------
I. Sched. Policy : basic_policy
E. Sched. Policy : N/A               E. Agg Rate Limit: 10000
I. Policer Ctrl. : N/A
E. Policer Ctrl. : N/A
Q Frame-Based Ac*: Disabled
Acct. Policy     : N/A               Collect Stats    : Disabled
Rad. Acct. Pol.  : N/A
Dupl. Acct. Pol. : N/A
ANCP Pol.        : N/A
HostTrk Pol.     : N/A
IGMP Policy      : N/A
MLD Policy       : N/A
Sub. MCAC Policy : N/A
NAT Policy       : N/A
Def. Encap Offset: none             Encap Offset Mode: none
Avg Frame Size   : N/A
Vol stats type   : full
Preference       : 5
Sub. ANCP-String : "iope-left-dupl"
Sub. Int Dest Id : "Gx-inserted-string"
Igmp Rate Adj    : N/A
RADIUS Rate-Limit: N/A
Oper-Rate-Limit  : 10000


show service id 10 subscriber-hosts detail
===============================================================
Subscriber Host table
===============================================================
Sap                    Subscriber
  IP Address
    MAC Address          PPPoE-SID Origin       Fwding State
---------------------------------------------------------------
[1/1/5:6.2]            iope-sub
  192.168.0.11
    00:00:65:06:02:01    N/A       DHCP         Fwding
---------------------------------------------------------------
Subscriber-interface  : int1
Group-interface       : g1
Sub Profile           : prem
SLA Profile           : voip+data
App Profile           : apps-prof
Egress Q-Group        : N/A
Egress Vport          : N/A
Acct-Session-Id       : D896FF0000001852EDFF46
Acct-Q-Inst-Session-Id: D896FF0000001952EDFF51
Address Origin        : DHCP
```

```
OT HTTP Rdr IP-FltrId : N/A
OT HTTP Rdr Status    : N/A
OT HTTP Rdr Fltr Src  : N/A
HTTP Rdr URL Override : N/A
GTP local break-out   : No
```

16.5.13 PCC Rules

A generic use case for flow-based dynamic policy is related to customized network level treatment of on-demand services. Such services can represent a wide range of applications, such as video-on-demand or access to a specific application in the network. The service can be identified by traffic destination parameters or DSCP bits. Once the service is identified, a set of actions can be applied to the service (rate change, forwarding-class change, Usage-Monitoring, and so on).

Typical flow of events for service activation is shown in Figure 227:

Figure 227 **Generic Use Case For IP Criterion-Based Policy Change via Gx**

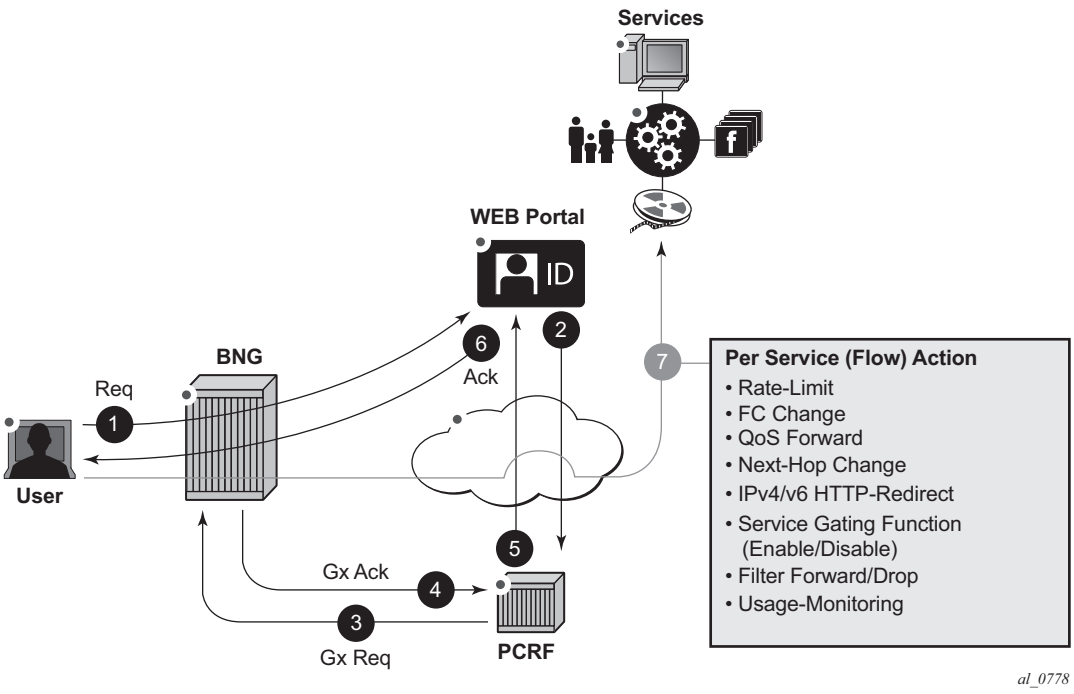

1) An established user subscribes to a service in the network via a Web portal at any given time.

2) Once the authentication/payment is accepted, the back end (for example, the Web portal integrated in OSS) identifies the service and submits the parameters defining the network delivery of the offered service to the PCRF.

3) The PCRF converts those parameters into rules and submits those rules to the subscriber-host on the BNG via the Gx. The rules identify the service on the network level (destination IP@ and port) along with the desired action.

4) [and 5) and 6)] Before the service can be started, the action of individual policy management elements must be acknowledged to ensure that the resources for the service delivery are available and instantiated before the service is delivered to the subscriber.

7) The service traffic can be started from the subscriber side. Network requirements for the successful service delivery are enforced on a per flow or DSCP basis as defined by the PCC rule.

16.5.13.1 PCC Rule Concept

A PCC rule consists of traffic classifiers (Flow-Information AVPs) required for traffic identification, and one or more actions associated with such classified traffic. PCC rules are unidirectional, which means that each rule is applied on ingress or egress. They are provisioned from PCRF via Gx interface.

Traffic classification is based on:

- 5 Tuple (IPv4 and IPv6)
- DSCP bits. When the content hosting device cannot be identified by the IP address, port and protocol, the DSCP marking can be used instead. Then, the DSCP marking is set by the client application and the markings should be preserved throughout the network until they reach the BNG.

Supported actions are:

- Rate limiting (in | out)
- Forwarding Class (FC) change (in | out)
- QoS forward (in | out)
- Next hop redirect (in)
- Service-Id redirect (in)
- HTTP redirect (in)
- Service gating function (in | out)
- Filter forward/drop (in | out)

- Usage-Monitoring (in | out)

16.5.13.2 PCC Rule Instantiation

A PCC rule that is submitted to the node via PCRF is internally instantiated using two basic policy constructs, QoS policy and filter policy (ACL). This internal division is transparent to the operator at the time of the rule provisioning. The operator perceives Gx as a unified method for provisioning policy rules, whether the rule is QoS-related or filter-related.

The type of action within the PCC rule will determine whether the PCC rule is split between the QoS policy and the filter policy.

Rules with actions:

- Rate limit
- Forwarding class change
- QoS forward — Matching traffic is forwarded without QoS actions and will not match on the next entry (match and exit behavior). This is equivalent to a *white-list* entry.
- Usage-Monitoring

are converted into a QoS policy while the rules with actions:

- Next hop redirect
- Service ID redirect
- HTTP redirect
- Gate Function (enable/drop)
- Filter forward/drop

are converted into **filter** rules.

The operator should be aware of this division for dimensioning (scaling) purposes. Operational commands can be utilized to reveal resource consumption on the node.

A PCC rule is addressed to a subscriber-host (single stack or dual stack) via the diameter session-id. However, qos-policy-related entries are applied per sla-profile instance since the qos resources are allocated per sla-instance. An sla-profile instance and sla-profile are two distinct concepts. An sla-profile instance is an instantiation of the sla-profile which is a configuration concept in which parameters are defined. An sla-profile is instantiated per a subscriber-host, or multiple subscriber-hosts can share an sla-profile instance as long as they belong to the same SAP and have the same subscriber ID.

This means that all hosts sharing the same sla-profile instance will inherit the change. The **sla-profile** *instance* and **sla-profile** are two distinct concepts. The **sla-profile** *instance* is an instantiation of the **sla-profile** which is a configuration concept in which parameters are defined. The **sla-profile** is instantiated per a subscriber-host, or multiple subscriber-hosts can share an **sla-profile** *instance* as long as they belong to the same SAP and have the same subscriber-id.

Filter-related entries are applied per each subscriber-host, whether the hosts are sharing or not sharing an sla-profile instance.

The concept of splitting the rules is shown in Figure 228.

Figure 228 **PCC Rule Conversion on the Node**

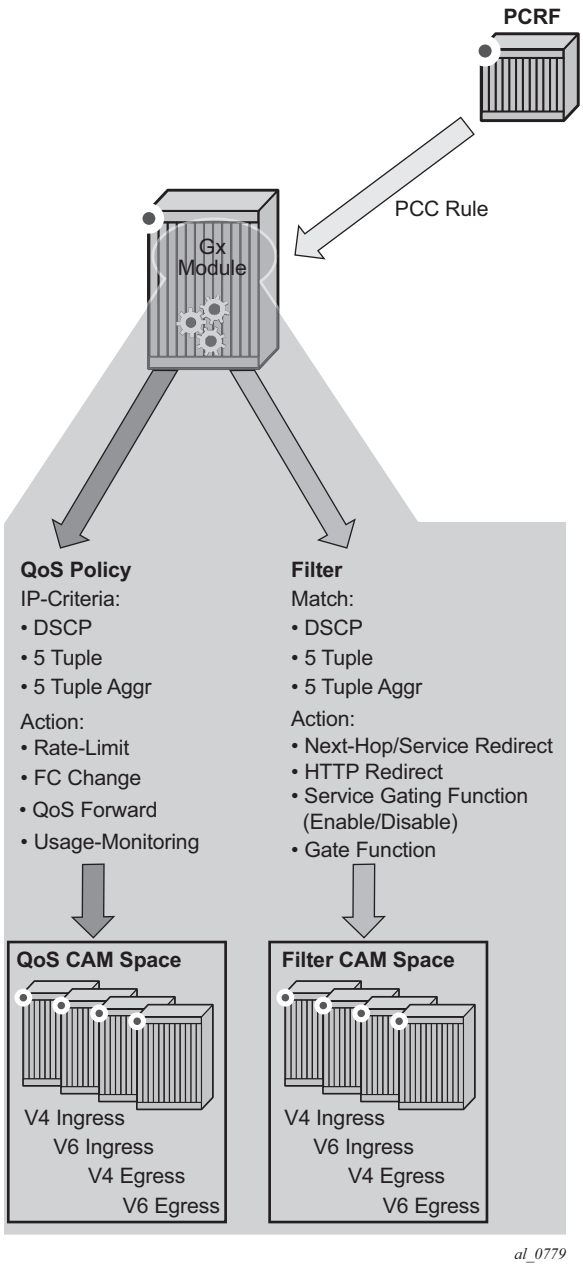

The PCC rule instantiation will fail if a PCC rule contains only actions without any classification, or if it contains only classification without any actions.

16.5.13.3 Base QoS-Policy and Base Filter

Subscriber host must have an explicit static (or base) filter or qos-policy before any dynamic entries can be inserted via Gx. For example, a base filter/qos-policy can be referenced by a sla-profile when the subscriber is instantiated. However, the parameters in the base qos-policy and base filter cannot be modified via Gx.

In the absence of explicitly defined qos-policy for the subscriber host, the default **qos-policy 1** will be in effect. Then, PCC rules with a QoS-related action cannot be applied.

PCC rule entries can be inserted in specifically allocated range in the base filter or qos-policy. The insertion point is controlled by the operator. This is shown in Figure 229. The entries reserved for PCC rules start at the beginning of the range specified by the following CLI command:

```
sub-insert-shared-pccrule start-entry <entry-id> count <count>
```

under the following CLI hierarchy:

```
configure>filter>ip-filter>
configure>filter>ipv6-filter
configure>qos>sap-ingress>
configure>qos>sap-egress>
```

An entry corresponds to a Flow-Information AVP and is equivalent to a match condition defined as any combination of the following parameters under a filter or **qos-policy ip-criteria**:

- source IP address
- destination IP address
- source port or port range
- destination port or port range
- protocol
- DSCP

Such defined entry will map into a single CAM entry with exception of port range configured as match criteria whereby a single port range command can expand into multiple CAM entries.

Static entries in filter/qos-policy can be inserted before and after the range reserved for PCC rules.

Figure 229 **Static and Dynamic QoS-Policy/Filter Entries**

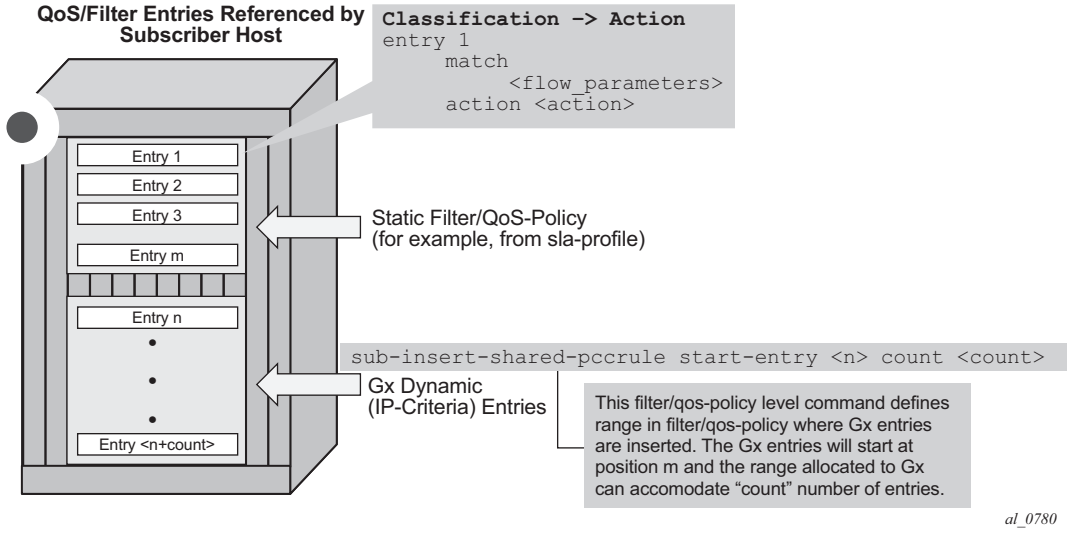

16.5.13.4 Generic Policy Sharing and Rule Sharing

Policy (defined in this context as a collection of static and dynamic rules) sharing between the subscriber hosts is depicted in Figure 230. In order to simplify CAM scaling explanations, the examples in this section assume that one rule within the policy occupies exactly one CAM entry. For simplicity, only PCC rules are shown but in reality a subscriber-host policy consist of PCC rules together with the base qos-policy/filter.

A policy, as a set of rules, can be shared amongst the subscriber-hosts. However, when a new rule is added to one of the subscriber-host, the newly created set of rules for this host becomes unique. Hence, a new policy for the subscriber-host is instantiated. This new policy consumes additional resources for all the old rules (clone of the old policy) along with the new rule. Figure below shows that a new policy (3) is instantiated when rule D is added to User 1, even though the rules A, B and C remain the same for Users 1 and 2. Policy 3 is a newly cloned with the same rules as Policy 1, and then Rule D is added onto it. On the other hand, when the rule C is applied to User 3, the set of rules becomes identical to the set of rules for User 2. Thus the two can start sharing rules and therefore the resources are freed.

Figure 230 **Policy Cloning**

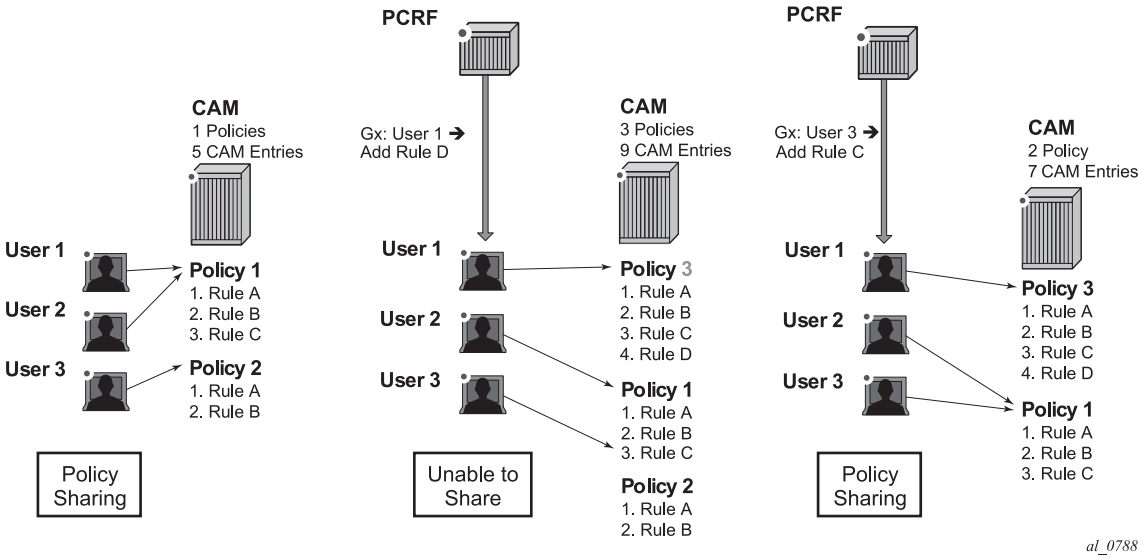

16.5.13.5 PCC Rule Name and PCC Rule Removal

Each PCC rule has a subscriber-host scope and it is referred to it by its name which is assigned by the operator on PCRF. The rules with exactly the same content but different rule names are evaluated into separate rules. To optimize performance and maximize scale, it is recommended that the rules sharing the same content have the same name (as provisioned in the PCRF).

PCC rules can be removed from the node via a Gx directive by referencing the PCC rule name. The rule name is supplied via the Charging-Rule-Name AVP at the time of the rule submission to the node by the PCRF. There is no Gx mechanism that would remove all PCC rules at once. Each PCC rule must be removed individually.

The AVP used to remove the rule from the node is:

```
Charging-Rule-Remove ::= < AVP Header: 1002 >
                                  *[ Charging-Rule-Name ]
```

An example of rule instantiation and rule removal is shown in Figure 231.

Figure 231 **Policy Removal by Name**

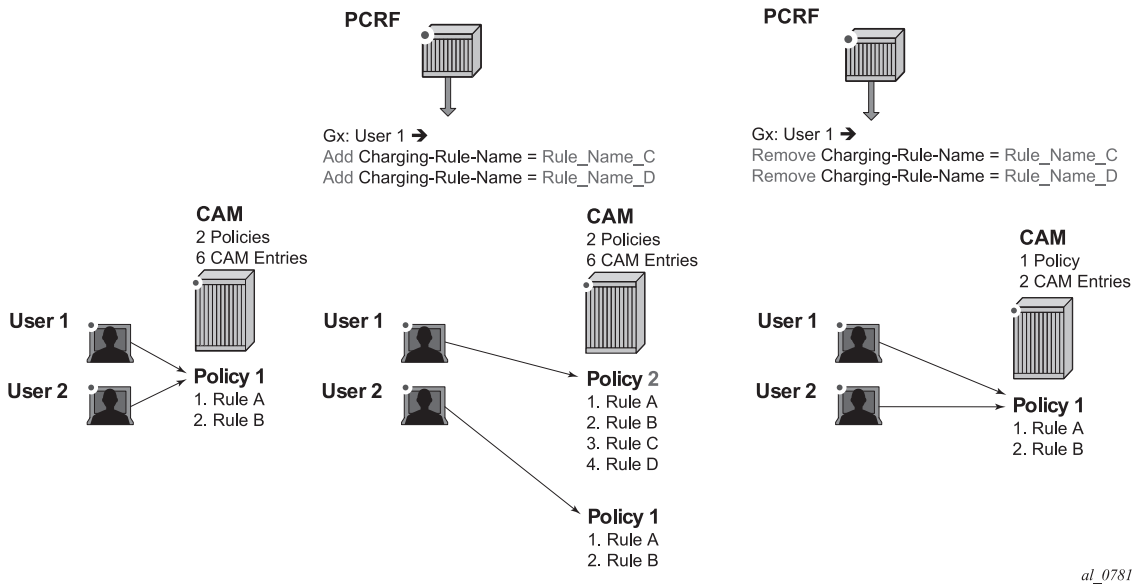

16.5.13.6 Gx Rule Ordering

Entries in IPv4/v6 filter and QoS policy created via CLI are ordered according to the numerical value associated with each entry command (which corresponds to the match condition) within the policy. CLI rules can be re-ordered with the **renum** command (in filters and QoS policies).

On the other hand, the PCC rules are ordered in one of the two ways. The difference between ordering of the entries with the rules, and the ordering of the rules themselves is:

- PCC rules are prioritized according to the Precedence AVP within the Charging-Rule-Definition AVP. They are inserted within the subscriber-host policy, according to the Precedence AVP relative to the other PCC rules already applied to the subscriber-host. A PCC rule with a lower Precedence value is applied before a rule with a higher Precedence value.

  The ordering behavior according to the Precedence is shown in Figure 232.

Figure 232 **PCC Rule Ordering - Priority Model**

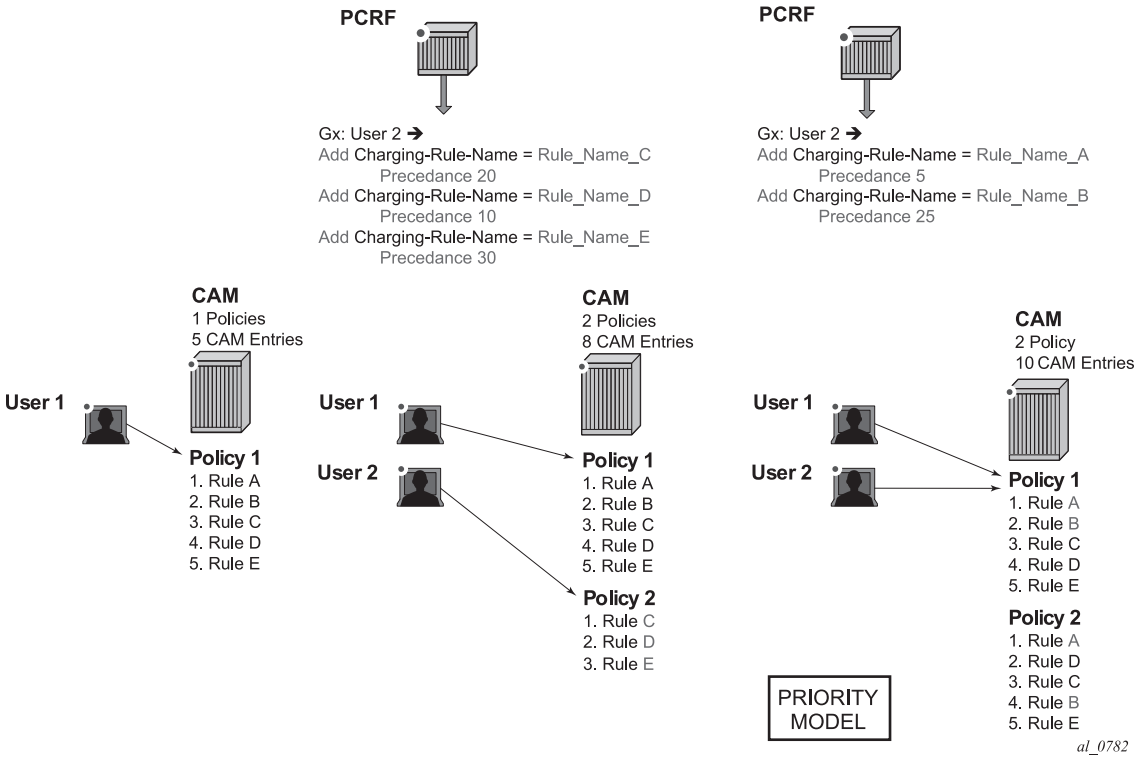

- Automatic optimization of PCC rules. Automatic optimization is used in cases where the PCC rule order is not important to the operator. Then, the node will optimize the rule ordering to achieve the best possible scale by maximizing rule sharing. This optimization (or internal rule ordering) is performed for PCC rules without the Precedence AVP, or for PCC rules with the same Precedence value.

  The premise of the automatic rule optimization is shown in Figure 232.

Figure 233 **Rule Ordering - Optimization Model**

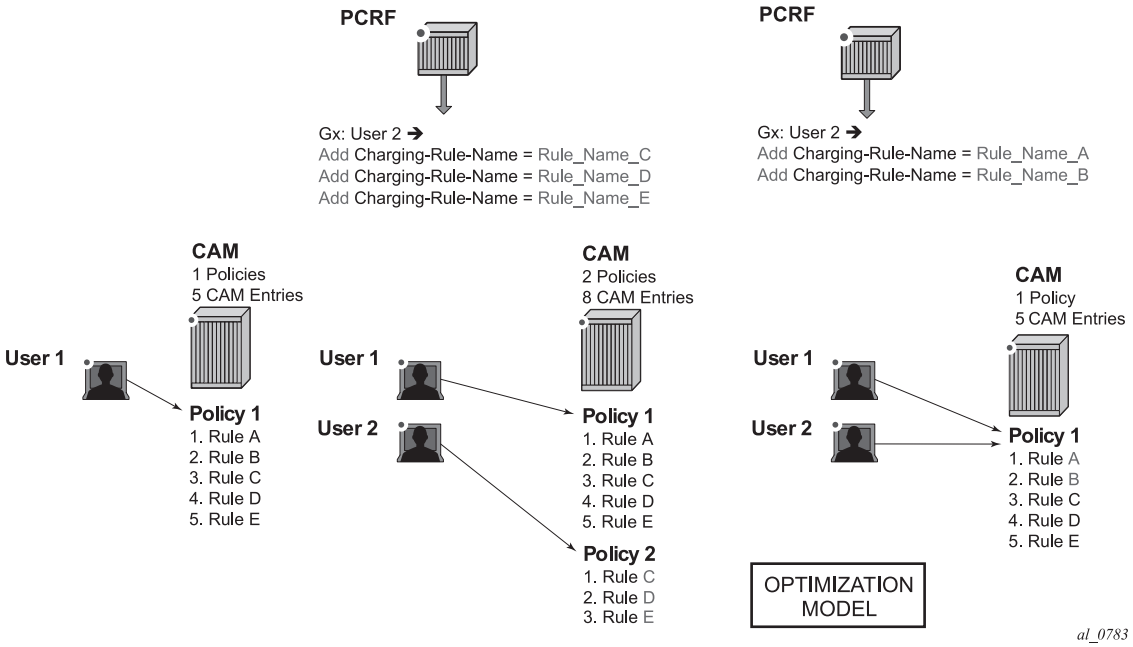

The ordering of PCC rules has no effect on the ordering of the static entries in the base qos-policy or filter.

Mixing of the PCC rules with the Precedence AVP and without Precedence AVP is allowed for the same subscriber-host. PCC rules without the Precedence AVP are inserted at the end of all PCC rules that do have the Precedence value set explicitly. In other words, the Precedence value for PCC rules without the explicitly configured Precedence AVP is assumed to be the highest. The PCC rules without the Precedence value are automatically inserted at the bottom of the PCC rule range.

A distinction should be made between the order of PCC rules in a PCC rule set and the order between the entries within each PCC rule. A PCC rule contains a group of classifiers that are all associated with the same actions. Therefore, the order of the entries (equivalent to match conditions) within any given PCC rule does not matter (all entries result in the same action). For this reason, PCC rules with identical name and identical entries but different order of the entries are automatically ordered in a way that would allow more optimal sharing of the rules between different subscribers.

16.5.13.7 PCC Rule Override

A PCC rule applied to a subscriber-host on the node can be overridden by re-submitting the PCC rule with the same name but different contents.

If at least 1 new flow is sent in the PCC rule update, then the existing flows are removed and replaced with the new flow. If no new flows are submitted, then the existing flows stay in place.

If there are conflicting parameters between the existing rule and the modified rule (for example the combination of the unsupported actions), the PCC rule override will fail.

16.5.13.8 Aggregation of IP-Criterion

An action with a PCC rule can be applied for a set of IP-criterion.

For example, a single policer can be instantiated for a set of flows for rate-limiting purposes.

A pseudo Gx directive would look like this:

```
Charging-Rule-Install — Directive to install the rule in 7750 SR
          Charging-Rule-Definition — PCC rule definition created on PCRF
               Charging-Rule-Name = Rule-1 — PCC rule name
                    Flow1     — match-criteria for flow 1
                    Flow2     — match-criteria for flow 2
                    Flow 2    — match-criteria for flow 3
                    Rate-limit    — rate-
limiting action applicable as an aggregate action for all 3 flows
```

All three flows are fed into the same rate limiter (policer).

16.5.13.9 Combining IPv4 and IPv6 Entries within the Rule

IPv4 flow entries and IPv6 flows entries can be combined within the same PCC rule. The actions that carry the IP address are address-type-specific (for example **next-hop-redirect**). All other actions (rate-limit, FC change, and so on) are universal and it are applied to both flow types (IPv4 and IPv6). The node will automatically sort out flow types (IPv4 and IPv6) within the rule and apply corresponding actions.

If the rule contains a mismatching flow type and actions (for example, IPv4 flows and IPv6 specific actions), the rule is rejected. It is the operator’s responsibility to ensure that the address-type-specific actions in the rule have corresponding flows to which they can be applied.

16.5.13.10 Gx Rules with Multiple Actions and Action Sharing

PCC rules can contain multiple actions. For the list of support action combinations, refer to the PCC Rules.

16.5.13.11 Alc-NAS-Filter-Rule-Shared AVP vs Flow-Information AVP

A Gx rule (as defined in a single Charging-Rule-Definition AVP) can contain either Flow-Information AVP or Alc-NAS-Filter-Rule-Shared AVP, but not both simultaneously.

Presence of either AVP within the Charging-Rule-Definition AVP determines the mode of operation for the rule:

Alc-NAS-Filter-Rule-Shared AVP indicates the mode of operation in which the permit or deny action is part of the flow definition itself (Alc-NAS-Filter-Rule-Shared AVP). This mode of operation is referred as NAS filter inserts. The basic format of the AVP is the following (RFC 4849 and 4005; AVP Code 400):

```
<action> <direction> <protocol>  from <source> to <destination> <options>.
```

There can be multiple ip-criteria definitions within the rule per subscriber-host, and each ip-criteria carries its own permit/deny action. There can be only one such rule (Charging-Rule-Definition) per subscriber-host. The rule entries are installed within the filter range defined by the following command:

```
sub-insert-shared-radius start-entry <entry-id> count <count>
under the following CLI hierarchy:
configure>filter>ip-filter>
configure>filter>ipv6-filter
```

Such rule cannot be removed by the Charging-Rule-Remove directive referencing the rule name. Instead, each such Gx rule will overwrite the previous one.

Flow-Information AVP indicates the mode of operation whereby all the flows in the rule share the same actions carried in separate AVPs. This mode of operation is referred to as PCC rule inserts. The rule entries are installed within the filter or qos-policy range defined by the following command:

```
sub-insert-shared-pccrule start-entry <entry-id> count <count>
under the following CLI hierarchy:
configure>filter>ip-filter>
configure>filter>ipv6-filter>
configure>qos>sap-ingress>
configure>qos>sap-egress>
```

There can be multiple flow based rules present in an orderly fashion and each rule can be individually removed by referencing its name.

Both modes of operation are supported simultaneously for the subscriber host.

16.5.13.12 RADIUS and Gx Interaction

Gx and RADIUS (CoA) policy management interfaces are simultaneously supported for the same subscriber-host.

RADIUS and Gx share the same entries for filter entry inserts (NAS-Filter-Rules and Alc-NAS-Filter-Rule-Shared) and therefore the most recent insert will override the previous one. Similar logic applies to subscriber-string overrides and QoS overrides, where the most recent source, overrides the previous one.

However, PCC rules (IP-criteria based Gx rules) are provisioned in a separate filter ‘entry’ space from RADIUS and Gx filter inserts and therefore the PCC rules and RADIUS/Gx based filter inserts can independently coexist.

Filter/QoS-policy entry order is shown in Figure 234. The order of configuration blocks (static, PCC rules or NAS filter inserts) is configurable. For example, an operator can specify that static filter entries are populated before PCC rules which are then populated before NAS filter inserts.

Figure 234 **CAM Table Population**

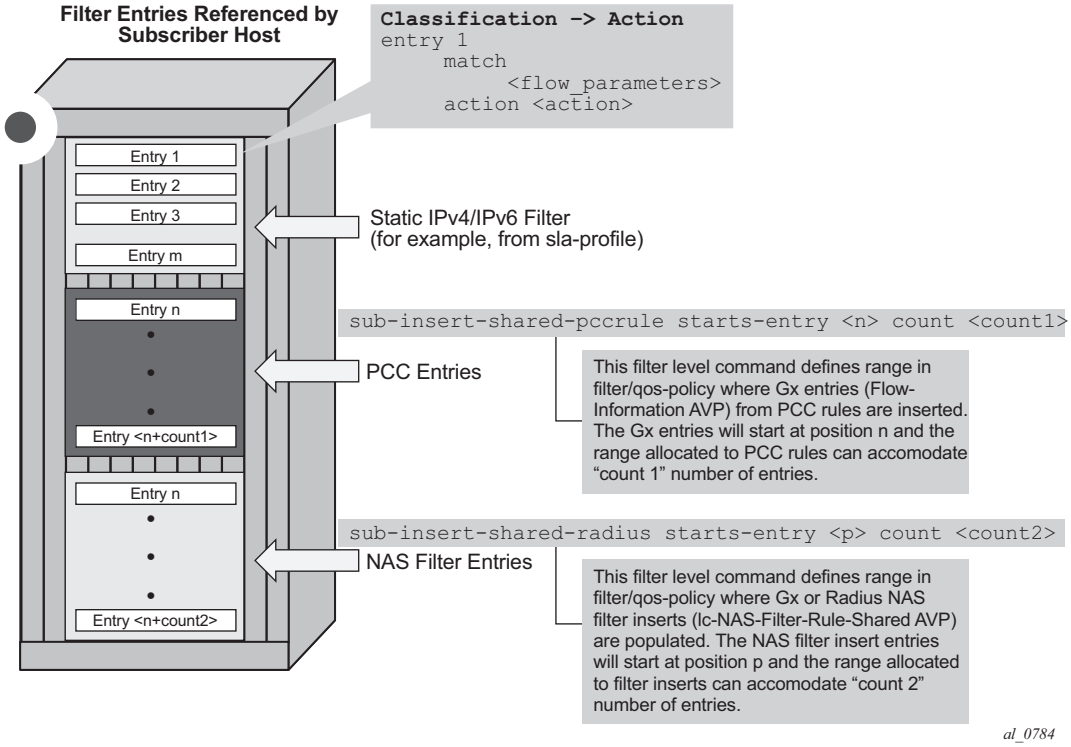

16.5.13.13 Bulk Changes via CLI while Gx Rules are Active

Once PCC rules are applied to a subscriber-host, the operator is allowed to modify via CLI some of the parameters in the base filter/qos-policy. For example, the operator is allowed to add/remove terms in the base ACL filter.

The list of the parameters in the QoS policy that can be changed is shown in Table 3. Adding/removing queue/policer, re-mapping of FC, modifying dscp-map or modifying static ip-criteria is not allowed.

Modified parameters in the base-policy/filter referenced in the sla-profile will affect all subscribers using this sla-profile. Replacing the base qos-policy/filter in sla-profile is not allowed for any subscriber-host if a clone of the base qos-policy/filter exist anywhere in the system.

However, replacing the base filter-id for a host via CoA or Gx override is allowed. Then, only the targeted host is affected and all existing PCC rules for this host are merged with the new filter.

Table 190 **CLI Modifiable Parameters in Base QoS Policy that Contains Clones**

CLI
config>qos>sap-ingress>queue
[no] cbs - Specify CBS
drop-tail low [no] percent-reduction-from-mbs - Specifies the drop tail for out-of-profile packets
[no] mbs - Specify MBS
[no] packet-byte-of* - Specify packet byte offset
[no] parent - Specify the scheduler to which this queue feeds
[no] percent-rate - Specify percent rates (CIR and PIR)
[no] rate - Specify rates (CIR and PIR)
config>qos>sap-egress>queue
[no] cbs - Specify CBS rate
drop-tail low [no] percent-reduction-from-mbs - Specifies the drop tail for out of profile packets
drop-tail exceed [no] percent-reduction-from-mbs - Specifies the drop tail for exceed profile packets
drop-tail high [no] percent-reduction-from-mbs - Specifies the drop tail for in profile packets
drop-tail highplus [no] percent-reduction-from-mbs - Specifies the drop tail for inplus profile packets
[no] mbs - Specify MBS rate
[no] parent - Specify the scheduler to which this queue feeds
[no] percent-rate - Specify percent rates (CIR and PIR)
[no] port-parent - Specify the port-scheduler to which this queue feeds
[no] rate - Specify rates (CIR and PIR)
[no] packet-byte-of* - Specify packet byte offset
config>qos>sap-ingress>policer
[no] cbs - Specify CBS
[no] high-prio-only - Specify high priority only percent-of-mbs
[no] mbs - Specify MBS
[no] packet-byte-of* - Specify packet byte offset
[no] parent - Specify the arbiter to which this policer feeds

Table 190 **CLI Modifiable Parameters in Base QoS Policy that Contains Clones (Continued)**

CLI
[no] percent-rate - Specify percent rates (CIR and PIR)
[no] rate - Specify rates (CIR and PIR)
config>qos>sap-egress>policer
[no] cbs - Specify Cbs
[no] high-prio-only - Specify high priority only percent-of-mbs
[no] mbs - Specify Mbs
[no] packet-byte-of* - Specify packet byte offset
[no] parent - Specify the scheduler to which this policer feeds
[no] percent-rate - Specify percent rates (CIR and PIR)
[no] rate - Specify rates (CIR and PIR)

16.5.13.14 PCC Rule Direction

PCC rules are unidirectional. The PCC rule direction is determined based on the value of the Flow-Direction AVP within the Flow-Information AVP. In the absence of the Flow-Direction AVP, the PCC rule direction is determined based on the Flow-Description AVP (as part of IPFilterRule direction field). Both of these AVPs (Flow-Direction and Flow-Description) are part of the PCC rule definition.

If the action within the PCC rule is in conflict with the direction of the flow, the PCC rule instantiation will fail. For example, an error is raised if the flow direction is UPSTREAM, while the action is ‘Max-Requested-Bandwidth-DL’ (downstream bandwidth limit).

16.5.13.15 Action

A PCC rule may contain multiple actions. Each action is carried in a separate, action specific AVP. The action specified in the **flow-description->ipfilterrule** data type is ignored. If rule contains multiple instances of the same action, each with a different value, the last occurrence of the action value is in effect.

Not all of the action types can be applied at the same time. The allowed combination of the actions per direction is given in Table 190 and Table 191.

16.5.13.16 Rate-Limiting Action (Ingress, Egress)

Rate-limiting action is implemented via policers. The policer is dynamically created at the PCC rule instantiation time. The rate can be enforced based on Layer 2 rates or Layer 3 rates.

Dynamically instantiated policers have their own policer id range to avoid the conflict with static policers.

The dynamically created policers will share common properties configured under the dynamic-policer CLI hierarchy:

```
configure
   qos
      sap-ingress/egress
         dynamic-policer
            stat-mode <stat_mode>
         parent <arbiter-name> [weight <weight-level>] [level <level>
         mbs <size> [bytes | kilobytes]
         cbs  <size> [bytes | kilobytes]
         packet-byte-offset {add <add-bytes> | subtract <sub-bytes>}
         range start-entry <entry-id> count <count>
```

The policer rates are part of PCC rule itself and are not part of static configuration.

The generic Gx directive for rate-limiting action is:

```
Charging-Rule-Install ::=      <AVP Header: 1001>
     Charging-Rule-Definition <AVP Header: 1003>
          Charging-Rule-Name <AVP Header: 1005>
          QoS-Information <AVP Header: 1016>
               Max-Requested-Bandwidth-
UL <AVP Header: 516> [bps] 3GPP 29.214 §5.3.15
         Max-Requested-Bandwidth-DL <AVP Header: 515> [bps] 3GPP 29.214 §5.3.14
         Guaranteed-Bitrate-UL <AVP Header: 1026> [bps] 3GPP 29.214 §5.3.26
         Guaranteed-Bitrate-DL <AVP Header: 1027> [bps] 3GPP 29.214 §5.3.25
```

The above rate limits refer to PIR and CIR rates of the dynamic policer in the respective direction.

16.5.13.16.1 Dynamic Policers and Queue Mappings

Once traffic is processed by the dynamic policers on **ingress**, the traffic will flow through the policer-output-queues shared queues. Traffic through dynamic policers will always bypass subscriber queues or policers on **ingress** that are statically configured in the base QoS policy.

Similar behavior is exhibited when static policers are configured on **egress**. Traffic outputting dynamic policer is never mapped to another static policer. Instead, such traffic is mapped to the corresponding shared queue in a queue-group. By default, this queue-group is the policer-output-queue group. However, the selection of the queue-group is configurable.

In contrast to the above, traffic processed by dynamic policers can be fed into statically configured **subscriber (local) queues on egress**. Dynamic policers and subscriber queues are tied through the forwarding-class.

The policer to local queue mapping and inheritance of the forwarding-class is shown in Figure 15. In this example, the mapping of traffic —> forwarding-class in rule 2 (flow 2) will depend on the DSCP bits in the traffic flow. If the DSCP value in this traffic flow are different from the explicitly configured DSCP values in the static (base) QoS policy, then traffic is mapped to the default forwarding-class.

Figure 235 **FC Inheritance for Dynamically Instantiated Policers on Egress**

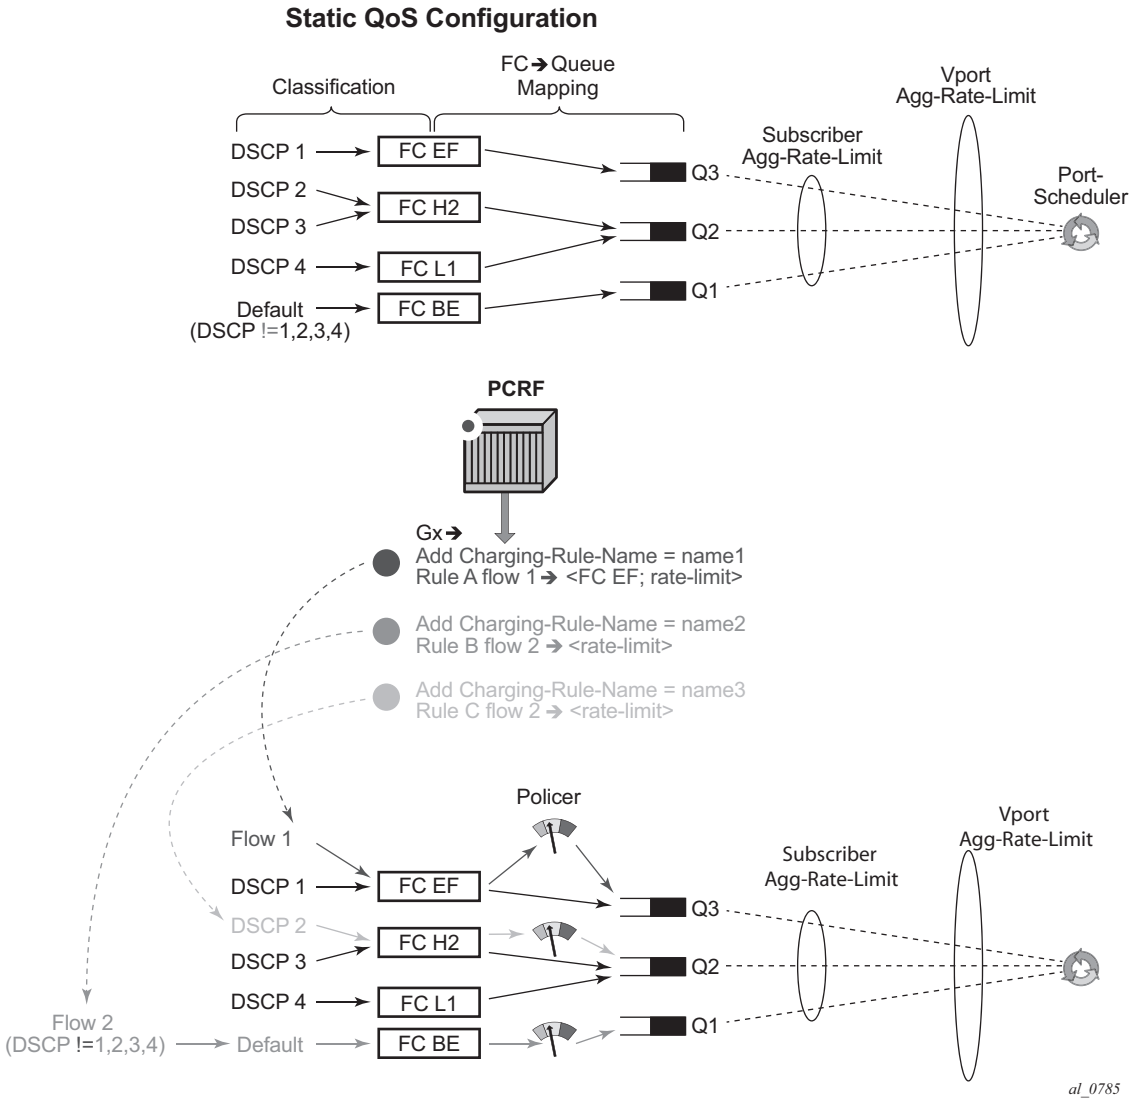

16.5.13.16.2 Dynamic Policer Rates and Accounting Statistics

By default, policer rates are configured based on Layer 2 frame length (for example, the Ethernet header plus the IP packet). This can be changed by the **packet-byte-offset** (PBO) command under the policer. If the policer is fed into a local queue, the PBO of the policer will not affect the PBO of the local queue it is feeding.

The rates for local subscriber queues can be independently measured based on Layer 2 or Layer 1 frame length and the queue statistics can be measured based on Layer 1, Layer 2 or Layer 3 (IP-only) frame length. The IP-only stats for queues can be configured in the **sub-profil>volume-stats-type {ip|default}.**

Dynamic policer (instantiated due to rate limiting or usage-monitoring action in PCC rules) statistics are not reported in RADIUS-based accounting. On egress, this will have no effect on volume counters in RADIUS-based accounting, since the dynamic policers are normally fed into local queues whose statistics are reported in RADIUS-based accounting. However, on ingress, the dynamic policers are always fed into the queue-group queues which are excluded from RADIUS based accounting. The consequence is that the ingress RADIUS-based accounting will lack statistics for the traffic that is flowing via dynamic policers.

If the dynamic policer is feeding a local queue, the aggregate statistics in show commands for such queue are not reported in order to avoid double counting (since the traffic statistics in **show** commands are already reported for the dynamic policer). However, the per-queue statistics are reported in **show** commands, irrespective of whether the policer is mapped to the local queue or not.

To avoid losing aggregate SAP or subscriber stats in show commands, the recommendation is to have policers feed into local queues which are not already mapped to an FC. For example:

```
queue 4 create //Not counted since policer 2 is feeding it
exit
policer 2 create
exit
fc be create
     queue 4     //Not counted
exit
fc l1 create
     queue 4     //Not counted
exit
fc ef create
     policer 2 queue 4
exit
```

FC BE, FC2 L1 —> queue 4

FC EF —> policer 2, queue 4

Then, traffic from queue 4 will not be counted in aggregate stats at all and consequently the aggregate accounting information is lost for FC BE and FC L1.

16.5.13.17 Forwarding-Class Change (Ingress, Egress)

Traffic can be re-prioritized via PCC rule by re-classification into a different forwarding class. The forwarding-class can be changed in several cases.

The original static mapping between traffic type, forwarding-class and the queue/ policer in the base qos-policy is configured outside of the ip-criteria CLI hierarchy.

For example:

```
configure>qos>sap-egress#
     dscp af11 fc "af"
```

Such mapping is configured outside of CAM and as such it has lower evaluation priority than the mapping configured via PCC rule which is installed in CAM.

The original static mapping is provisioned in the base qos-policy via ip-criteria CLI hierarchy.

For example:

```
config>qos>sap-egress>#
      ip-criteria
           entry 40000 create
                match
                     dscp af11
            exit
       action fc "af"
  exit
 exit
```

Then, the configured entry range for PCC rules **must** precede the static entry (match criteria) in which the original forwarding-class is configured. The insertion point (entry) is controlled via configuration: **sub-insert-shared-pccrule start-entry <***entry-id***> count <***count***>** command under the qos-policy.

In both of the above cases, the following PCC rule would override forwarding-class for traffic with DSCP value of 10 (af11 traffic class) from value af to h2.

```
Charging-Rule-Install
     Charging-Rule-Definition
          Charging-Rule-Name = fc-change
          Flow-Information
               ToS-Traffic-Class = 00101000 11111100
               Flow-Direction = 1
          QoS-Information
               QoS-Class-Identifier  = 2
```

The eight forwarding classes are mapped to QCIs (3GPP TS 23.203 §6.1.7.2) in the following manner:

BE —> QCI 8

L2 —> QCI 7

AF —> QCI 6

L1 —> QCI 4

H2 —> QCI 2

EF —> QCI 3

H1 —> QCI 1

NC —> QCI 5

The generic Gx directive for forwarding-class change:

```
Charging-Rule-Install ::=      <AVP Header: 1001>
     Charging-Rule-Definition <AVP Header: 1003>
          Charging-Rule-Name <AVP Header: 1005>
               QoS-Information <AVP Header: 1016>              3GPP 29.212 §5.3.16
                    QoS-Class-Identifier <AVP Header: 1028>     3GPP 29.212 §5.3.17
```

16.5.13.18 QoS Forward (Ingress and Egress)

Create an **ip-criteria** and/or **ipv6-criteria** entry with no action specified. Matching traffic is forwarded without a QoS action and will not match on a next entry (match and exit behavior). This is equivalent to a *white-list* entry.

CLI equivalent:.

```
config>qos
  sap-ingress | sap-egress <id> create
    ip-criteria | ipv6-criteria
      entry <id> create
        match
          <5-tuple | dscp>
        exit
        action
      exit
    exit
```

The generic Gx directive for QoS forwarding:

```
Charging-Rule-Install ::= <AVP Header: 1001>
```

```
Charging-Rule-Definition <AVP Header: 1003>
    Charging-Rule-Name <AVP Header: 1005>
        Alc-QoS-Action = Forward (1) <AVP Header: 1028>
```

16.5.13.19 Next Hop Redirect (Ingress)

The next hop redirection will explicitly or implicitly change the next hop for the traffic flow within the same service ID (routing context) or a different service ID (routing context).

If the next hop is not explicitly provided, the next hop is selected automatically, according to the routing lookup in the referenced service ID.

The generic Gx directive:

```
Charging-Rule-Install <AVP Header: 1001>
     Charging-Rule-Definition <AVP Header: 1003>
          Charging-Rule-Name <AVP Header: 1005>
               Alc-Next-Hop :: <AVP Header: 1023>
                    Alc-Next-Hop-IP <AVP Header: 1024>
Alc-V4-Next-Hop-Service-id <AVP Header: 1025>
                   Alc-V6-Next-Hop-Service-id <AVP Header: 1026>
```

This action overwrites the routing table lookup based on the destination IP and sets the next hop to the:

- IPv4/6 address within the same service ID
- IPv4/6 address within a different service ID

The next hop search is indirect, which means that if the explicitly provided next hop in the PCC rule cannot be found in the routing table, then an additional routing table lookup is performed to find the path (next hop) to the indirect next hop from the PCC rule.

If only the service-id is specified in PCC rule (without the next hop), then the next hop is selected from the specified service-id based on the destination IP address of the packet.

16.5.13.20 HTTP Redirect (Ingress)

HTTP redirect utilizes Redirect-Information AVP from 3GPP 29.212, §5.3.82.

The generic Gx directive:

```
Redirect-Information < AVP Header: 1085 >
     Redirect-Support < AVP Header: 1086 >
     Redirect-Address-Type < AVP Header: 433 >
     Redirect-Server-Address < AVP Header: 435 >
```

16.5.13.21 Filter Forward/Drop (Ingress and Egress)

This action is used to control traffic flow within a PCC rule by using ALU specific AVP. PCC rules are utilizing filters and QoS policies as distinct building blocks. This action within a PCC rule will create an IP or IPv6 filter entry with an action **forward** or **drop**.

The CLI equivalent follows:

```
config>filter
    ip-filter | ipv6-filter <id> create
        entry <id> create
            match
                <5-tuple | dscp>
            exit
            action
                forward | drop
            exit
        exit
    exit
```

The generic Gx directive for filter forward/drop is implemented through ALU specific AVP:

```
Charging-Rule-Install ::= <AVP Header: 1001>
    Charging-Rule-Definition <AVP Header: 1003>
        Charging-Rule-Name <AVP Header: 1005>
            Alc-Filter-Action = Forward (1) <AVP Header: 1027>
            Alc-Filter-Action = Drop (2)
```

16.5.13.22 Service Gating Function

The service gating function is used to enable or disable the service that is represented by the PCC rule. This action is enforced through a Flow-Status AVP (AVP code 511) - 3GPP 29.214, §5.3.11.

The system supports the following values (actions) for the Flow-Status:

- enabled (2) — Enables all actions specified within the PCC rule.

  Note that although Flow-Status (service gating function) is considered an action, in this context it is used to enable all other actions that are explicitly set within the PCC rule.

- disabled (2) — Disable all action specified within the PCC rule and drops the traffic.

The service gating function is applicable in the direction that is associated with the rule (PCC rules in the system are unidirectional).

Flow-Status is by default enabled (2) (if the Flow-Status AVP is not explicitly specified within the PCC rule). Flow-Status=Enabled must be accompanied by one or more additional actions in the same PCC rule (refer to Gx Rules with Multiple Actions and Action Sharing for a list of allowed simultaneous actions), otherwise the PCC rule instantiation in the node will fail.

If the Flow-Status is set to ‘disabled (3)’, all other actions within the same rule loses their meaning since the packet is dropped. Basically, the ‘disabled’ directive disables the flow of packet through the system. A disabled Flow-Status is equivalent to the Alc-Filter-Action = Drop (2).

This AVP is carried inside of Charging-Rule-Definition (3GPP 29.212, §5.3.5):

```
Charging-Rule-Definition ::= < AVP Header: 1003 >
                              { Charging-Rule-Name }
                             *[ Flow-Information ]
                              [ Flow-Status ]
                              [ QoS-Information ]
                              [ Precedence ]
                             *[ Flows ]
                              [ Monitoring-Key]
                              [ Redirect-Information ]
                             *[ AVP ]
```

16.5.13.23 PCC Rule Provisioning Example

The following is an example of PCC rule provisioning in a CCA-I message:

```
<CC-Answer> ::=  < Diameter Header: 272, PXY >
     < Session-Id >
     { Auth-Application-Id }
     { Origin-Host }
     { Origin-Realm }
     [ Result-Code ]
     [ Experimental-Result ]
     { CC-Request-Type }
     { CC-Request-Number }
     *[ Supported-Features ]
     *[ Event-Trigger ]
     [ Origin-State-Id ]

     Charging-Rule-Install ::= <AVP Header: 1001> -> host instantiation
          Charging-Rule-Name = “ingr-v4:7”
          Charging-Rule-Name = “eggr-v6:5”
          Charging-Rule-Name = “Sub-Profile:prem”
```

```
          Charging-Rule-Name = "Sla-Profile:voip+data"
          Charging-Rule-Name = "Inter-Dest:vport-AN-1"

     Charging-Rule-Install  -> service instantiation
          Charging-Rule-Definition
               Charging-Rule-Name = "service-1"  -
> should be able to remove the rule by name later on
                    Flow-Information      -> traffic flow definition
                    Flow-Description = "permit in 6(TCP) from any to ip 10.10.10.10/
32 40000-40010"
                    ToS-Traffic-Class = 00101000 11111100]  -
> DSCP definition (value mask). In case of the DSCP, Flow-
Direction (1080) AVP must be included.
                    Flow-Direction = UPSTREAM  -> traffic flow direction
                    QoS-Information <AVP Header: 1016>
                         Max-Requested-Bandwidth-UL = 10000000  -
> UPSTREAM rate definition (not downstream, since the traffic flow direction is IN)
                         QoS-Class-Identifier  = 3      -
> EF forwarding class; in general one of 8 forwarding-
classes (FC) in 7450 ESS and 7750 SR (be|l2|af|l1|h2|ef|h1|nc). This is used for re-
prioritization of the traffic.
```

In this example the host is instantiated using the two Charging-Rule-Install AVPs. The first one is used to instantiate the host. The second one is used to instantiate the IP-criterion based service named service-1. Service-1 is defined as the upstream traffic flow with traffic class AF11, destined to the TCP port range 40000-40010 on the node with IP address 10.10.10.10/32.

The actions for this traffic flow are:

- Rate-limit of 10M
- Change forwarding-class to AF11.

16.5.13.24 Operational Aspects

The commands used to examine dynamic rules and NAS filter inserts associated with the subscriber hosts are shown in Figure 236.

Figure 236 **Overview of PCC Rule-Related Operational Commands**

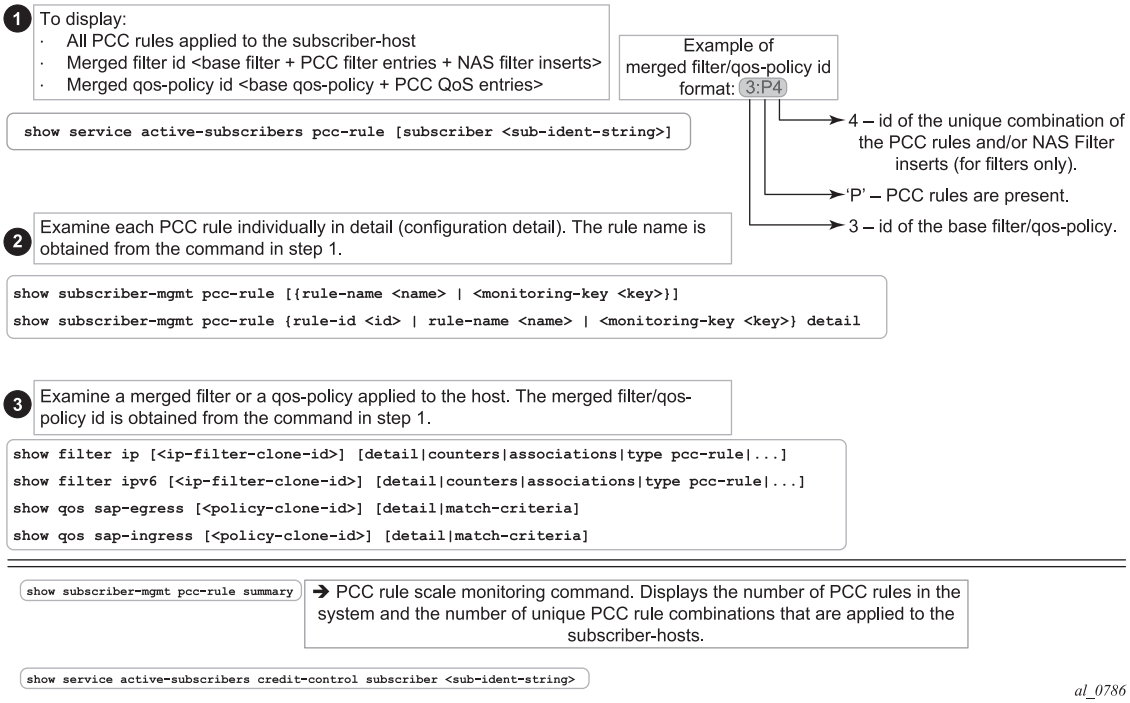

16.5.13.25 PCC Rules and Capacity Planning

One of the most important factors to be considered for capacity planning with PCC rules is the number of unique policies that are applied to subscribers.

A unique policy constitutes a base a QoS policy or filter ID along with all PCC rules that are applied to a subscriber or a set of subscribers.

Now examine an example where there are ‘n’ PCC rules in the system (‘n’ qos rules and ‘n’ ACL filter rules). Those rules are applied to IPv4 traffic in ingress direction. Further, assume that the PCC rules do not have defined Precedence AVP, which means that the system can optimize their order for maximum sharing and maximized scale. Then, ‘n’ PCC rules can be combined by various permutation into 2^n-1 unique combinations Next assumption is that there are five possible base qos-policies for IPv4 traffic in ingress direction and five possible base filters for IPv4 traffic in the ingress direction.

Given the above, the unique PCC rule combinations (2^n-1) together with five base qos-polices will produce 5*(2^n-1) unique qos-policies per ingress IPv4. Same logic can be applied for ingress IPv4 filters.

This exercise must be repeated for egress direction as well as for IPv6 type traffic, by taking into consideration the number of respective base qos-policies/filters and the number of PCC rules.

Once the number of unique policy combinations is determined and ensured that it is within the system limits, each policy must be further evaluated to determine the number of entries it will take in CAM.

16.5.13.26 PCC Rule Scaling Example

Figure 237 depicts an example relevant to capacity planning with focusing on understanding the scaling limits pertaining to the number of PCC rules and their mutual combinations when they are applied to the subscriber hosts.

This example is focuses on an IPv4 filter applied in ingress direction but similar logic can be used in understanding other policy types (QoS, egress, IPv6).

The system/line card limits in this example are set to the following values for illustration purposes only:

- Number of PCC rules per system: 1K
- Max number of rules per subscriber host: 64
- Max number of combinations of the PCC rules (or services) that are active (applied to the subscriber-hosts) per system: 4K
- Max ingress IPv4 filter CAM entries per FP2/FP3: 64k
- Max filters per system: 16k

Note: The actual CAM limits vary per policy type (filter/QoS), direction and IP address type (v4 vs v6). The actual scaling limits can be found in the Scaling Guides for the relevant software release.

Figure 237 **Example of the Scaling Limits for PCC Rules**

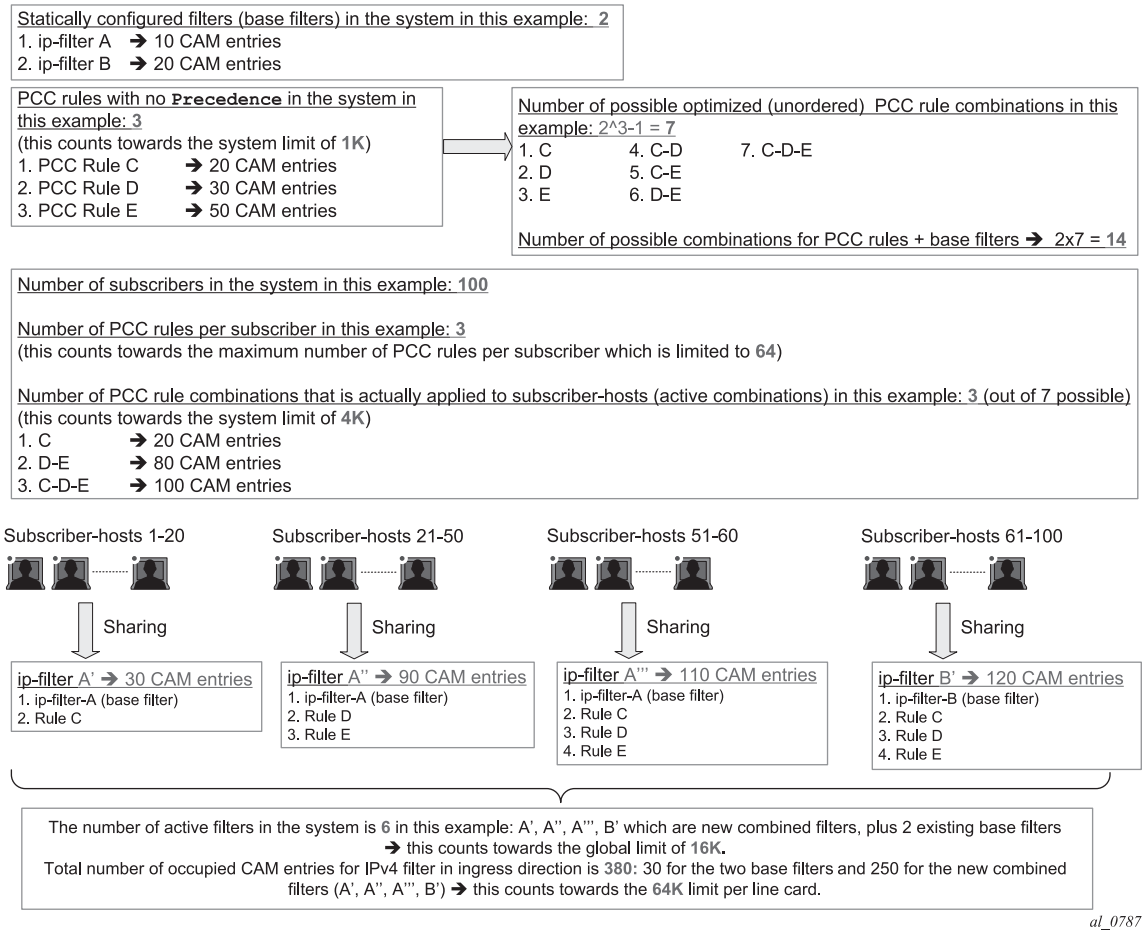

16.5.14 NAS Filter Inserts

Gx filter entries inserted via the NAS-Filter-Rule are subscriber-host-specific entries. This means that in the upstream direction, the source IP address in the NAS-Filter-Rule will always be internally set by the node to the IP address of the subscriber host itself. Similarly, in the downstream direction, the destination IP address in the NAS-Filter-Rule is set by the node to be the IP address of the subscriber-host itself.

On the other hand, the entries in the Alc-NAS-Filter-Rule-**Shared** AVP are processed as received without any modifications. This means that such entries can be shared with all the hosts that have the same Alc-NAS-Filter-Rule-Shared applied.

Similar to QoS overrides, NAS filter entries are not predefined on the node but instead they are defined under the Charging-Rule-Install — Charging-Rule-Definition AVP.

The Charging-Rule-Name AVP for NAS filter inserts is an arbitrary name that is part of Charging-Rule-Definition AVP in which NAS-Filter-Rule AVP or Alc-NAS-Filter-Rule-Shared is provided. Such Charging-Rule-Name is used to report errors related to instantiation of the inserts.

16.5.14.1 Examples of NAS Entry Inserts

The following AVPs identify NAS filter inserts that are applied to a subscriber host. Those AVPs can be included in CCA-i, CCA-u or RAR message sent from the PCRF.

```
Charging-Rule-Install ::= <AVP Header: 1001>
Charging-Rule-Definition <AVP Header: 1003>
            Charging-Rule-Name <AVP Header: 1005> = “allow-all”
            Alc-NAS-Filter-Rule-
Shared <AVP Header: 158> = "permit in ip from any to any "ASCII NUL" permit out ip
from any to any"
```

In this example, the filter entry defined in Alc-NAS-Filter-Rule-Shared AVP is inserted in the clone of the existing base filter for the subscriber(s).

16.5.15 Error Handing and Rule Failure Reporting in ESM

The Gx rule (overrides, PCC rules or NAS filter inserts) instantiation failure can occur on two levels:

- AVP Decoding level in the Diameter — The Gx message contains an unrecognized AVP with the M-bit set. Then, all Gx rules (ESM, UM and AA) in the message are rejected and the CCR-u with the Charging-Rule-Report AVP (rule status) and Error-Massage AVP (failure description) or an RAA message with the appropriate Result-Code AVP (fail – 5xxx), Error-Message AVP (description) and the Failed-AVP AVP are sent to the PCRF.

Invalid content of a supported AVP with the M-bit set will also trigger the same response. Invalid content of an AVP refers to the malformed syntax of an AVP that carries the type of the AVP value implicitly embedded in the AVP value itself. Consider sla-profile:rule-name-1 string. then, the sla-profile: refers to the type of the value carried in the Charging-Rule-Name AVP. The value that the Charging-Rule-Name carries is rule-name-1 and this value represents the sla-profile name already configured on the node (as opposed to filter, sub-profile, category-map, and so on). If the sla-profile is misspelled (type is unrecognized), the whole AVP is un-decodable.

- Gx rule instantiation level in ESM, UM or AA – Then each module (ESM, UM or AA) would fail all rules destined for it. The failure of a Gx rule within a module can be caused by referencing non-existing profile (for example sla-profile:unknown-name) or a lack of resources on the node. Then, a CCR-u message from the respective module is sent with the Rule-Report-Status AVP listing all the rules destined for this module and the corresponding Error-Message AVP describing the cause for the failure.

  Another example that can cause all rules with the ESM module to fail would be invalid combination of actions within the rule.

16.5.15.1 AVP Decoding Failure in Gx

Reporting an AVP decoding problem in Gx is described in the following example:

A Gx directive is received to install two overrides on the node. The two overrides are supposed to change the sla-profiles and sub-profiles for the subscriber host. The AVP that is used to change the sla-profile is miss-formatted. The predefined sla-profile keyword in the Charging-Rule-Install AVP is misspelled as spa-profile instead of sla-profile.

```
Charging-Rule-Install ::= <AVP Header: 1001>
    Charging-Rule-Name <AVP Header: 1005> = "Sub-Profile:prem"
    Charging-Rule-Name <AVP Header: 1005> = "Spa-Profile:voip+data"
```

Since the Charging-Rule-Name AVP has the M-bit set, the whole message will fail and an error is reported. No rules within this Gx message is installed (not even the valid ones, then this would be the Charging-Rule-Name = "**Sub-Profile**:prem").

Note: If the M-bit was clear in the Charging-Rule-Name AVP, the erroneous AVP would be simply ignored and proceed with installation of the remaining, correctly formatted rules.

The nature of the error depends on the original directive sent by the PCRF (RAR or CCA – push or pull model).

If the directive from the PCRF is passed with the **cca** command, the response is CCR-u with the following error related AVPs:

```
[ Error-Message ] — "Invalid value spa-profile:voip+data"

Charging-Rule-Report ::= < AVP Header: 1018 >
    *[ Charging-Rule-Name ] — Spa-Profile:voip+data
    [ PCC-Rule-Status ]    — INACTIVE (1)
    [ Rule-Failure-Code ]  — GW/PCEF_MALFUNCTION (4)


Charging-Rule-Report ::= < AVP Header: 1018 >
    *[ Charging-Rule-Name ]      —  Sub-Profile:prem
    [ PCC-Rule-Status ]          —  INACTIVE (1)
    [ Rule-Failure-Code ]        —  GW/PCEF_MALFUNCTION (4)

Failed-AVP ::= < AVP Header: 279 >
    Charging-Rule-Name = Spa-Profile:voip+data
```

- If the directive is passed to the node through RAR, the node responds with the following RAA message:

```
Failed-AVP ::= < AVP Header: 279 >
    Charging-Rule-Name = Spa-Profile:voip+data

Result-Code ::= < AVP Header: 268 > = DIAMETER_INVALID_AVP_VALUE (5004)
```

Similarly, if the number of filter entries for each entry type (NAS-Filter-Rule — host-specific or Alc-NAS-Filter-Rule-Shared — shared) exceeds the maximum supported number (see the *7750 SR Gx AVPs Reference Guide*), the whole message will fail the decoding phase.

The reason that the Result-Code AVP is present in the RAA message and not in the CCR-u message is that this code is only allowed to be present in the answer messages, according to the standard.

16.5.15.2 ESM Rule-Installation Failure

This assumes that the rule installation directives are successfully passed from the Gx module to the ESM module and the failure to install rules occurs in the ESM module.

In the Gx override example below, the referenced sla-profile is unknown. Then, all directives passed to the ESM module will fail and consequently no rules or overrides are installed. The sub-profile change fails as well although the prem sub-profile is known in the system.

```
Charging-Rule-Install ::= <AVP Header: 1001>
    Charging-Rule-Name <AVP Header: 1005> = "Sub-Profile:prem"
            Charging-Rule-Name <AVP Header: 1005> = "Sla-Profile:unknown"
```

The error reporting flow is as follows:

- If the directives are passed using the CCA command, the response is CCR-u command with the following error related AVPs:

```
[ Error-Message ]                    – "sla-profile 'unknown' lookup failed"

Charging-Rule-Report ::= < AVP Header: 1018 >
    *[ Charging-Rule-Name ]                –  Sla-Profile:unknown
    [ PCC-Rule-Status ]                    –  INACTIVE (1)
    [ Rule-Failure-Code ]                  –  GW/PCEF_MALFUNCTION (4)


Charging-Rule-Report ::= < AVP Header: 1018 >
    *[ Charging-Rule-Name ]          –  Sub-Profile:prem
    [ PCC-Rule-Status ]              –  INACTIVE (1)
    [ Rule-Failure-Code ]            –  GW/PCEF_MALFUNCTION (4)
```

- If the directive is passed to the node with RAR, the node responds with the following messages:

  RAA = OK since the Gx module successfully processed the AVP parsing.

  The RAA is followed by CCR-u, triggered by the rule instantiation failure in ESM module. CCR-u will contain the following AVP related to the rule status:

```
[ Error-Message ]                –   "sla-profile 'unknown' lookup failed"

Charging-Rule-Report ::= < AVP Header: 1018 >
    *[ Charging-Rule-Name ]       –  Sla-Profile:unknown
    [ PCC-Rule-Status ]           –  INACTIVE (1)
    [ Rule-Failure-Code ]         –  GW/PCEF_MALFUNCTION (4)


Charging-Rule-Report ::= < AVP Header: 1018 >
    *[ Charging-Rule-Name ]       –  Sub-Profile:prem
    [ PCC-Rule-Status ]           –  INACTIVE (1)
    [ Rule-Failure-Code ]         –  GW/PCEF_MALFUNCTION (4)
```

Similar behavior would be exhibited if the directive is sent to the UM or AA modules. However, ESM, UM and AA are separate modules and failure to install rule in one module will not affect rule installation in another.

16.5.15.3 Failure Reporting in AA

Failure reporting in AA is performed in similar fashion as in ESM.

Instead of Charging-Rule-Report AVP, the ADC-Rule-Report is used:

```
ADC-Rule-Report ::= < AVP Header: 1097 >
                        *[ ADC-Rule-Name ]
                         [ PCC-Rule-Status ]
                         [ Rule-Failure-Code ]
                        *[ AVP ]
```

16.5.15.4 Summary of Failure Reporting

Table 191 summarizes Gx failure reporting on the node.

Table 191 **Failure Reporting**

Failure Event	Gx Message Received via CCA (Pull Model)	Gx Message Received via RAR (Push Model)
AVP decoding/ interpreting failure; M-bit cleared	Ignore AVP	Ignore AVP
AVP decoding/ interpreting failure; M-bit set	CCR-u is sent by the node. CCR-u will contain: • Charging-Rule-Report AVP for all rules (all rules inactive) • First failed AVP in Failed-AVP AVP • Error-Message AVP at the top level describing the reason for the failure. No rules within the message is instantiated on the node.	RAA is sent by the node. RAA will contain: • Result-Code AVP [DIAMETER_INVALUID_AVP_ VALUE (5004), DIAMETER_AVP_UNSUPPORTED (5001), DIAMETR_UNAMBLE_TO_COMPLY (5012) • First failed AVP in Failed-AVP AVP No rules within the message is instantiated on the node.

Table 191 **Failure Reporting (Continued)**

Failure Event	Gx Message Received via CCA (Pull Model)	Gx Message Received via RAR (Push Model)
Rule failure in ESM	CCR-U is sent by the node. CCR-u will contain: • Charging-Rule-Report AVP for all rules (all rules inactive) • Error-Message AVP at the top level describing the reason for the failure. No rules is instantiated in the ESM module.	RAA with the Result-Code AVP 'success' (2001) is sent by the node, followed by a CCR-u. CCR-u will contain: • Charging-Rule-Report AVP for all rules (all rules inactive) • Error-Message AVP at the top level describing the reason for the failure. No rules is instantiated in the ESM module.
Rule failure in Usage-Monitoring (UM)	CCR-U is sent by the node. CCR-u will contain: • Charging-Rule-Report AVP for all rules (all rules inactive) • Error-Message AVP at the top level describing the reason for the failure. No rules is instantiated in the UM module.	RAA with the Result-Code AVP 'success' (2001) is sent by the node, followed by a CCR-u. CCR-u will contain: • Charging-Rule-Report AVP for all rules (all rules inactive) • Error-Message AVP at the top level describing the reason for the failure. No rules is instantiated in the UM module.
Rule failure in AA	CCR-U is sent by node. CCR-u will contain: • ADC-Rule-Report AVP for all rules (all rules inactive) • Error-Message AVP at the top level describing the reason for the failure. No AA rules is instantiated in the AA module.	RAA with the Result-Code AVP 'success' (2001) is sent by the node, followed by CCR-u. CCR-u will contain: • ADC-Rule-Report AVP for all rules (all rules inactive) • Error-Message AVP at the top level describing the reason for the failure. No rules is instantiated in the AA module.

16.5.16 Usage-Monitoring and Reporting

Usage-Monitoring and reporting refers to the collection and reporting of octets (volume) that a service or application on the node has consumed during a certain period. The usage on the node is reported via the Gx interface to the PCRF. Based on this information, the PCRF can apply a specific action (policy change) to the entity being monitored. For example, QoS can be modified, or the service can be blocked when specific thresholds are reached.

Usage-Monitoring and reporting is performed over a single Gx session for the ESM/AA subscriber. In other words, there is only a single session for an ESM subscriber-host and corresponding AA subscriber. Via this single Gx session, Usage-Monitoring can be requested simultaneously in ESM context (PCC rule level, credit-category and/or IP-CAN session) and AA context (application based Usage-Monitoring).

16.5.16.1 ESM Usage-Monitoring - What is Being Monitored

In the ESM context, volume consumption (octets - 3GPP 23.203 §4.4) can be monitored on three levels:

- per entire IP-CAN session
- per credit-category
- per PCC rule

Usage-Monitoring can be monitored simultaneously on all three levels.

An IP-CAN session on the node represents a subscriber-host whose service types are determined by the sla-profile instance. In per IP-CAN session volume monitoring, the aggregated queue or policer counters are reported per direction (in or out). This includes dynamic policers that are instantiated as a result of a Gx action; for example, rate-limiting.

The following configuration is necessary to allow per IP-CAN session level Usage-Monitoring to be enabled for sessions associated with the category map:

```
configure
    subscriber-mgmt
        category-map <category-map-name> [create]
            gx-session-level-usage
```

If the sla-profile instance changes mid-session, the counters are reset.

One obvious difference between regular RADIUS accounting and Gx Usage-Monitoring is that in RADIUS accounting the cumulative byte number for sla-profile instance is presented in each report (interim-updates or stop acct messages), while in Usage-Monitoring this count is reset between the two reports (when the quota is reached, the usage report is triggered).

Per credit-category monitoring refers to volume monitoring of a single queue/policer or a set of queues/policers within the sla-profile instance. Each queue/policer (or set of queues/policers as a subset of the sla-profile instance) represents a service for which the Usage-Monitoring is required. Those queues/policers (services) are organized on the node in credit categories.

```
*A:7750>config>subscr-mgmt>cat-map# info
----------------------------------------------
            activity-threshold 1
            credit-exhaust-threshold 50
            category "queue1" create
                queue 1 ingress-egress
            exit
            category "queue3-5" create
                queue 3 ingress-egress
                queue 5 ingress-egress
            exit
            category "rest-queues" create
                queue 2 egress-only
                queue 4 egress-only
                queue 6 egress-only
                queue 7 egress-only
                queue 8 egress-only
            exit
----------------------------------------------
```

Each service category has a name that is used to reference the category in Usage-Monitoring and reporting.

The category-map (predefined on the node) that is used in Usage-Monitoring can be associated with the subscriber-host through the following methods (in the order of priority):

- PCRF — Charging-Rule-Install AVP that references the category map in Charging-Rule-Name = cat-map:<cat-map-name>
- LUDB
- RADIUS
- Python script

PCC rule Usage-Monitoring reports volume usage per flow or set of flows. PCC rule Usage-Monitoring is described in a separate section below.

Usage-Monitoring for the subscriber host can be configured on the node, but it will not be active until it is turned on by the PCRF either via CCA-i, CCA-u or RAR.

Usage-Monitoring can be enabled per ingress and/or egress direction or as total count. However monitoring the total count is mutually exclusive with per direction count. For example, total Usage-Monitoring cannot be enabled simultaneously with ingress (or egress) Usage-Monitoring for the same monitoring entity (session or category).

16.5.16.2 AA Usage-Monitoring – What is Being Monitored

In AA, charging groups (CG), application groups (AG) and applications are monitored. Refer to the *7450 ESS, 7750 SR, and VSR Multiservice Integrated Service Adapter Guide* for details.

16.5.16.3 Requesting Usage-Monitoring in ESM

Gx Usage-Monitoring is activated explicitly from the PCRF via CCA-I, CCA-U or RAR. It is triggered via the Usage-Monitoring-Information AVP along with the event-trigger = usage-report (33). The Usage-Monitoring-Information AVP contains the following AVPs:

```
Usage-Monitoring-Information::= < AVP Header: 1067 >
                                      [ Monitoring-Key ]
                               0,2    [ Granted-Service-Unit ]
                               0,2    [ Used-Service-Unit ]
                                      [ Usage-Monitoring-Level ]
                                      [ Usage-Monitoring-Report ]
                                      [ Usage-Monitoring-Support ]
```

There could be multiple instances of Usage-Monitoring-Information AVP present in a single CCA or RAR messages. For example, simultaneous Usage-Monitoring for IP-CAN session level, credit-category level or pcc rule level can be requested.

Usage-Monitoring-Level for IP-CAN session is set to SESSION_LEVEL (0)

Usage-Monitoring-Level for category-map is set to PCC_RULE_LEVEL (1)

Usage-Monitoring-Level for PCC rules is set to PCC_RULE_LEVEL (1)

16.5.16.4 Reporting Accumulated Usage

The node reports usage information to the PCRF under the following conditions:

- When a usage threshold is reached
- When all pcc rules associated with the monitoring are removed or deactivated
- When Usage-Monitoring is explicitly disabled by PCRF
- When a session is terminated
- When requested by PCRF (on demand)

To report accumulated usage for a specific monitoring-key, the node sends a CCR with the Usage-Monitoring-Information AVP containing the accumulated usage information since the last report. For each of the enabled monitoring-keys, the Usage-Monitoring-Information AVP will include the Monitoring-Key AVP and the accumulated volume usage in the Used-Service-Unit AVP.

A usage report on the node can be triggered by reaching the usage threshold communicated to the node by the PCRF in the CCR-u message carrying accumulated usage for that monitoring entity along with the Event-Trigger AVP set to USAGE_REPORT.

In response to the CCR-u message, the PCRF will communicate to the node via a CCA-u message whether the Usage-Monitoring should continue:

- If the new thresholds for the currently monitored entity/levels are provided in Granted-Service-Units AVP, the Usage-Monitoring will continue
- If the thresholds are not included in Granted-Service-Units AVP, the Usage-Monitoring will stop.

Thresholds are incremental. For example, if the quota of 100 MB is submitted to the node, the usage should be reported when that quota is reached. At that point, the user can be granted another 100 MB. The new usage report on the node is triggered when another 100 MB are accumulated. Absence of the threshold for a given entity in the CCA-u message is an indication that the Usage-Monitoring should stop.

When the PCRF informs the node that Usage-Monitoring should stop (by not including thresholds in CCA-u), the node will not report usage which has accumulated between sending the CCR and receiving the CCA.

Another possibility of usage reporting is on-demand. In this scenario, usage for one or more monitoring keys is reported whether the usage threshold has been reached. This is achieved by sending the node the Usage-Monitoring-Report AVP (within the Usage-Monitoring-Information AVP) set to USAGE-MONITORING_REPORT_REQUIRED. If the Monitoring-Key AVP is omitted in such a request, Usage-Monitoring for all enabled entities is reported to the PCRF.

If that the credit-category is removed from the subscriber host (the sla-profile instance referencing the category-map is changed for the subscriber host), the node will report the outstanding usage in a CCR-u message with the Event-Trigger set to USAGE_REPORT.

16.5.16.5 Disabling Usage-Monitoring

When the PCRF explicitly disables Usage-Monitoring on the node, the node will report the accumulated usage which has occurred while Usage-Monitoring was enabled.

To disable Usage-Monitoring for an entity, the PCRF sends the Usage-Monitoring-Information AVP referencing only the applicable monitoring entity with the Monitoring-Key AVP and the Usage-Monitoring-Support AVP set to USAGE_MONITORING_DISABLED.

When the PCRF disables Usage-Monitoring in a RAR or CCA command, the node sends new a CCR-U with the Event-Trigger AVP set to "USAGE_REPORT" to report the accumulated usage for the disabled Usage-Monitoring entities.

16.5.16.6 Usage-Monitoring for PCC Rules

Each PCC rule for which Usage-Monitoring is required, contains Monitoring-Key AVP.

```
Charging-Rule-Definition ::= < AVP Header: 1003 >
                              { Charging-Rule-Name }
                             *[ Flow-Information ]
                              [ Flow-Status ]
                              [ QoS-Information ]
                              [ Precedence ]
                              [ Monitoring-Key]
                             *[ AVP ]
```

Usage-Monitoring for PCC rules is implemented through a dynamic policer. The policer is instantiated at the time when the PCC rule with Monitoring-Key AVP is installed.

The same monitoring-key can be used in multiple PCC rules assuming that these rules are for the same direction. In other words, the charging rule is rejected if the same monitoring-key is used for ingress and egress.

16.5.16.7 Session Termination

At IP-CAN session termination, the node sends the accumulated usage information for all entities for which Usage-Monitoring is enabled in the CCR-t.

16.5.16.8 Usage Monitoring When Multiple Subscriber Hosts or Sessions Share an SLA Profile Instance

A Diameter Gx session from which Usage Monitoring is started controls the Usage Monitoring for the entire SLA profile instance. Only one Diameter Gx session can control the Usage Monitoring per SPI at a given time. That is, Usage Monitoring can only be started from a single GX session when multiple subscriber hosts or sessions share an SLA profile instance.

The session ID of the Diameter Gx session that controls the Usage Monitoring is displayed as the “Diameter Session Gx” field in the output of the following **show** command.

```
# show service active-subscribers credit-control [subscriber <sub-ident-string>]
===============================================================================
Active Subscribers
===============================================================================
-------------------------------------------------------------------------------
Subscriber sub-1 (sub-profile-1)
-------------------------------------------------------------------------------
-------------------------------------------------------------------------------
(1) SLA Profile Instance sap:1/1/2:10 - sla:sla-profile-1
-------------------------------------------------------------------------------
Category Map         : catmap-1
Diameter Session Gx  : bng1.domain.com;1490015297;7
Number of categories
static               : 2
gx-session           : 0
gx-pcc               : 1
Category Name        : cat-1
Ingress Queues       : 1
Egress Queues        : 1
Ingress Policers     : (Not Specified)
Egress Policers      : (Not Specified)
No monitoring
HTTP Rdr URL Override: (Not Specified)
Category Name        : cat-2
Ingress Queues       : 2
Egress Queues        : 2
Ingress Policers     : (Not Specified)
Egress Policers      : (Not Specified)
No monitoring
HTTP Rdr URL Override: (Not Specified)
Category Name        : fl_um1_up
Volume Used          : 0                    Volume Available     : 100000
HTTP Rdr URL Override: (Not Specified)
```

A Diameter Gx session stops being the controlling Gx session for Usage Monitoring of the SLA profile instance in the following situations:

- the Gx session controlling Usage Monitoring for an SLA profile instance is terminated, for example, because the associated subscriber session is disconnected. A CCR-T is sent immediately reporting the usage.
- all Usage Monitoring is disabled form the controlling Gx session. A CCR-U is sent immediately reporting the usage.

When the Usage Monitoring terminates, the usage of subscriber hosts or sessions sharing the SLA profile instance is no longer accounted for. Usage Monitoring can now be started from another Diameter Gx session, which then becomes the controlling Gx session for Usage Monitoring on the SLA profile instance.

16.5.16.9 Usage-Monitoring Examples

For the description of the specific AVP, refer to the *7750 SR Gx AVPs Reference Guide*.

IP-CAN session Usage-Monitoring

A **category-map** with **gx-session-level-usage-monitoring** must be associated with the subscriber host or session:

```
configure
    subscriber-mgmt
        category-map cat-map-1 create
            gx-session-level-usage
```

PCRF in RAR sends the following AVPs (among all the other mandatory ones: session-id, etc.)

```
Usage-Monitoring-Information
        Monitoring-Key = "any-string"
        Granted-Service-Unit
            CC-Input-Octets = 1000000
              CC-Output-Octets = 1000000
        Usage-Monitoring-Level = session_level(0)

Event-Trigger = USAGE_REPORT
```

The node reports usage when the thresholds are reached some time later in the CCR-U. The usage is monitored internally on the node based on the current sla-profile instance.

```
Usage-Monitoring-Information
        Monitoring-Key = "any-string"
        Used-Service-Unit
        CC-Input-Octets = 1000000
                 CC-Output-Octets = 1000000
```

The PCRF instructs the node to continue Usage-Monitoring with the new thresholds in the CCA-U:

```
Usage-Monitoring-Information
        Monitoring-Key = "any-string"
        Granted-Service-Unit
            CC-Input-Octets = 1000000
            CC-Output-Octets = 1000000
        Usage-Monitoring-Level = session_level(0)
```

Category Usage-Monitoring

Assume that the following category-map is associated with the subscriber host:

```
*A:7750>config>subscr-mgmt>cat-map# info
----------------------------------------------
            activity-threshold 1
            credit-exhaust-threshold 50
            category "queue1" create
                queue 1 ingress-egress
            exit
            category "queue3-5" create
                queue 3 ingress-egress
                queue 5 ingress-egress
            exit
            category "rest-queues" create
                queue 2 egress-only
                queue 4 egress-only
                queue 6 egress-only
                queue 7 egress-only
                queue 8 egress-only
            exit
```

The PCRF will send the following AVPs in the RAR message (among all the other mandatory ones: session-id, etc.)

```
Charging-Rule-Install
        Charging-Rule-Name = Cat-Map:cat1 – cat-map rule install

Usage-Monitoring-Information
        Monitoring-Key = "queue-1"
        Granted-Service-Unit
            CC-Input-Octets = 1000000
            CC-Output-Octets = 1000000
            Usage-Monitoring-Level = PCC_RULE_LEVEL (1)
Usage-Monitoring-Information
```

```
        Monitoring-Key = "queue-3-5"
        Granted-Service-Unit
            CC-Input-Octets = 2000000
            CC-Output-Octets = 2000000
            Usage-Monitoring-Level = PCC_RULE_LEVEL (1)
Event-Trigger = USAGE_REPORT
```

The node reports usage when the thresholds are reached some time later in the CCR-U:

```
Usage-Monitoring-Information
        Monitoring-Key = "queue-1"
        Used-Service-Unit
               CC-Input-Octets = 1000000
               CC-Output-Octets = 1000000
Usage-Monitoring-Information
        Monitoring-Key = "queue-3-5"
        Used-Service-Unit
               CC-Input-Octets = 2000000
               CC-Output-Octets = 2000000
```

The PCRF instructs the node to continue Usage-Monitoring with the new thresholds in the CCA-U:

```
Usage-Monitoring-Information
        Monitoring-Key = "queue-1"
        Granted-Service-Unit
            CC-Input-Octets = 1000000
            CC-Output-Octets = 1000000
            Usage-Monitoring-Level = PCC_RULE_LEVEL (1)
Usage-Monitoring-Information
        Monitoring-Key = "queue-3-5"
        Granted-Service-Unit
            CC-Input-Octets = 2000000
            CC-Output-Octets = 2000000
            Usage-Monitoring-Level = PCC_RULE_LEVEL (1)
```

16.5.17 Event Triggers

The PCRF may subscribe to an event trigger on the node. The PCRF subscribes to new event triggers or removes armed event triggers unsolicited at any time. When an event matching the event trigger occurs, the node reports the event to the PCRF. The event triggers that are required in procedures are unconditionally reported (for example, IP address allocation/de-allocation) from the node, while the PCRF may subscribe to the remaining events (for example Usage-Monitoring).

When sent from the PCRF to the node, the Event Trigger AVP indicates an event that will trigger an action in the node. When sent from the node to the PCRF, the Event Trigger AVP indicates that the corresponding event has occurred. If no Event Trigger AVP is included in a CCA or RAR operation, any previously provisioned event trigger will still be applicable.

The PCRF may remove all previously provided event triggers by providing the Event-Trigger AVP set to the value NO_EVENT_TRIGGERS. When an Event-Trigger AVP is provided with this value, no other Event-Trigger AVP is provided in the **cca** or **rar** command. Upon reception of an Event-Trigger AVP with this value, the node will not inform the PCRF of any event except for those events that are always reported and do not require provisioning from the PCRF.

When the PCRF subscribes to one or more event triggers by using the **rar** command, the node will send the corresponding currently applicable values to the PCRF in the RAA if available, and in this case, the Event-Trigger AVPs will not be included.

For a list of the supported events on the node, refer to the *7750 SR Gx AVPs Reference Guide*.

16.5.18 Subscriber Verification

At any time, the PCRF can query the node for the presence of the subscriber-host via a RAR message.

The node will respond with the following result-codes in RAA:

- DIAMETER_SUCCES (2001) — subscriber active
- DIAMETER_UNKNOWN_SESSION_ID (5002) — subscriber does not exist

16.5.19 Subscriber Termination

The PCRF can request IP-CAN session termination on the node via two messages:

- a RAR directive with the Session-Release-Cause AVP to the node
- ASR

Upon the arrival of either of those messages, the node starts the IP-CAN session termination procedure (CCR-t with a corresponding Termination-Cause AVP is sent to the PCRF). This is described in the 3GPP 29.212 document, §4.5.9.

For a list of the supported Termination-Cause AVP values on the node, refer to the *7750 SR Gx AVPs Reference Guide*.

16.5.20 Mobility Support in WiFi

When a WiFi subscriber moves between the access points (APs), a CCR-u message is triggered on the node, carrying the Called-Station-Id AVP. The Called-Station-Id AVP carries the MAC IP address of the new AP. This functionality allows the PCRF to make a location-based policy decision.

This functionality is enabled via event trigger USER_LOCATION_CHANGE (13) [3GPP 29.212, §5.3.7] sent to the node by a PCRF in a CCA or RAR message.

The same event is reported back from the node to the PCRF in a CCR-u message when the user location changes.

16.5.20.1 Redundancy

Redundancy in Gx relies on the Diameter redundancy mechanisms described in Diameter Redundancy.

16.5.21 Persistency and Origin-State-ID AVP

Persistency and Origin-State-ID AVP (RFC 6733, §8.6 and §8.16).

Persistency (saving the state of IPoE hosts on the compact flash) for Gx sessions is not supported. This means that, on reboot, the node will restore the subscriber-hosts from the persistency but the Gx session awareness for the recovered hosts is lost. Any previously applied QoS or filter overrides are lost. However, subscriber-strings (subscriber-id, sub-profile, sla-profile, aa-profile) can be made persistent and can be preserved across reboots.

The Origin-State-Id (OSI) AVP is not stored in persistency. If the node reboots, the Origin-State-ID AVP is set to boot time (UTC).

The Origin-State-Id AVP is contained in the CER messages and application messages that are sent from the node to the PCRF/DRA. In the other direction, sent by the PCRF to the node, the OSI is ignored.

To restore a lost session after the reboot, the node will initiate a CCR-i message for every host that has recovered from persistency. The CCR-i will contain the new session-id and origin-state-id. Based on this CCR-i, it is expected that the PCRF will return the most current policy for the host.

16.5.22 Overload Protection

Each SR OS node has a receiving queue per Gx application (ESM, UM, AA). Each queue can hold 10,000 messages. While the queue is in the overloaded state, the SR OS node will reply to every new RAR message with the RAA (ACK) immediately followed by a CCR-U message containing the error-message with the description ‘Overload’. This can be considered as explicit signaling towards the PCRF notifying it of the condition on the SR OS node.

If the messages in the overwhelmed queue do not require sending an answer (in case that the overwhelmed queue contains CCA-I/U messages), the TCP window will fill up, TCP ACKs will not be sent and consequently this IS an implicit notification to the PCRF to slow down.

If the SR OS node receives a response from an overloaded PCRF (Result-Code = DIAMETER_TOO_BUSY), the SR OS node timeouts (tx-timer) the originally sent message. Once the message is timed out, the configuration settings (on-failure) determines whether to trigger the peer-failover procedure or not (Peer-failover based on DIAMETER_TOO_BUSY Result-Code is recommended in RFC6733, §7.1.3.

16.6 Diameter NASREQ Application

The Diameter NASREQ application is used for Authentication, Authorization, and Accounting services in the Network Access Server (NAS) environment. The SR OS supports a stateless operation of NASREQ authentication and authorization, interacting with a NASREQ server that does not maintain session state.

Subscriber host or session authentication results in an AA-Request (AAR) message being sent to the Diameter NASREQ server. An Auth-Session-State AVP with value equal to 1 (No State Maintained) is included in the AAR to inform the server of the stateless mode. The server responds with an AA-Answer (AAA) message and must include the Auth-Session-State AVP with value equal to 1 (No State Maintained), together with the authorization AVPs.

Diameter NASREQ accounting is not supported.

Table 192 lists the supported Diameter NASREQ messages. Vendor-specific AVP's are shown as: v-<vendor-id>-<AVP id>.

Table 192 **Supported Diameter NASREQ Messages**

Diameter Message		Code
AAR	AA-Request	265
AAA	AA-Answer	265

Diameter NASREQ authentication is supported for IPoE hosts and sessions, PPPoE PTA PAP/CHAP authentication. Diameter NASREQ authentication is not supported for L2TP LAC/LNS.

NASREQ and RADIUS authentication cannot be configured simultaneously on a capture-sap, local-user-database, or group-interface. They have the same priority in the hierarchy of different sources (such as local user database, Gx, defaults, etc.) for obtaining the subscriber host or session authorization parameters.

Multi-chassis redundancy is supported via separate Diameter NASREQ peers on each redundant node. Each node of the multi-chassis redundancy pair has its own Diameter Identity (origin host/realm). The subscriber host or session is authenticated on the BNG where it is initially connected. Due to the stateless operation, there is no need to synchronize NASREQ session state. Alternatively, the Diameter proxy can be used if it is required to have a single Diameter Identity (origin host/realm) per pair of multi-chassis redundant nodes.

There is no NASREQ re-authentication for active subscriber hosts or sessions, except for a forced re-authentication when the circuit ID/interface ID or remote ID of a DHCP host is changed.

Stateless NASREQ authentication can be complemented with Diameter Gx policy management for policy control and mid-session changes. Diameter NASREQ and Gx applications are supported simultaneously on a single Diameter peer.

Figure 238 shows a sample call flow for a subscriber using Diameter NASREQ for authentication and Diameter Gx for policy management.

Figure 238 **Sample Diameter NASREQ Call Flow**

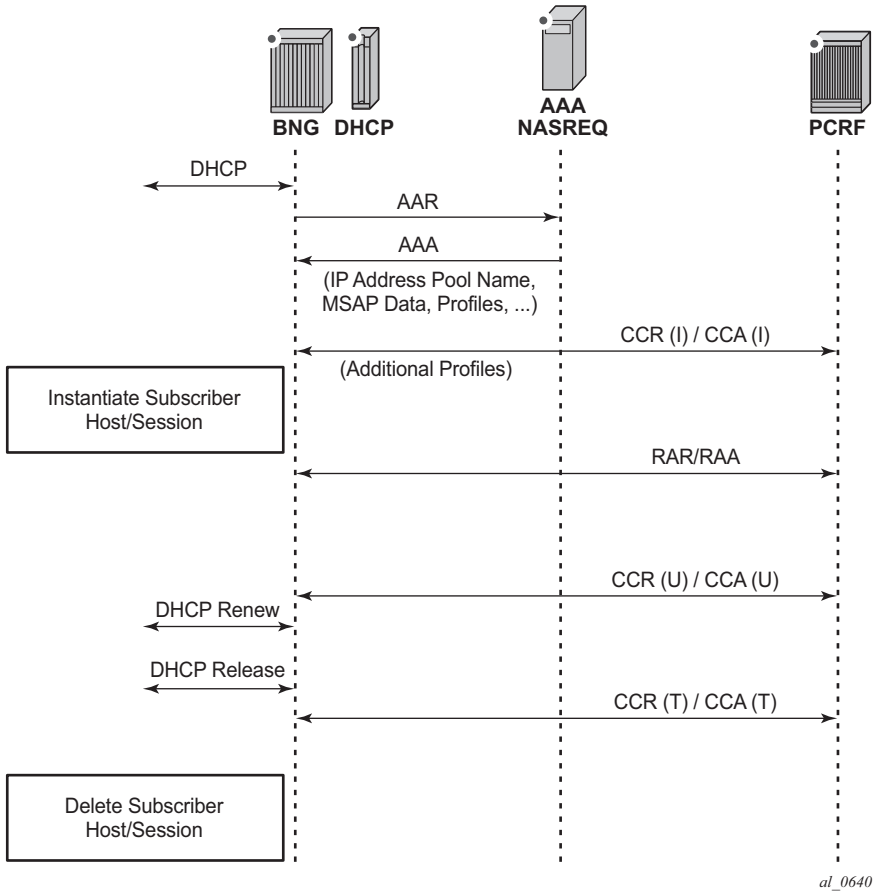

Table 193 lists the authorization AVPs that are accepted in a Diameter NASREQ AA-Answer message. Vendor-specific AVPs are shown in the table as: v-<vendor-id>-<AVP-id>.

Table 193 **AA-Answer Message — Accepted Authorization AVPs**

AVP ID	AVP Name	Description
1	User-Name	Overrides the “Radius User-Name”.
8	Framed-IP-Address	The IPv4 address of the subscriber host.
9	Framed-IP-Netmask	The IPv4 netmask of the subscriber host.
22	Framed-Route	IPv4 managed route to be configured on the NAS for a routed subscriber host.
25	Class	Opaque value; echoed in RADIUS accounting.
88	Framed-Pool	The name of an IPv4 address pool.
97	Framed-IPv6-Prefix	SLAAC IPv6 prefix (wan-host).
99	Framed-IPv6-Route	IPv6 managed route to be configured on the NAS for a v6 routed subscriber host.
100	Framed-IPv6-Pool	The name of an IPv6 IA-NA address pool (wan-host).
123	Delegated-IPv6-Prefix	DHCPv6 IA-PD IPv6 prefix (pd-host).
26.6527.9	Alc-Primary-Dns	The IPv4 address of the primary DNS server.
26.6527.10	Alc-Secondary-Dns	The IPv4 address of the secondary DNS server.
26.6527.11	Alc-Subsc-ID-Str	Unique subscriber ID string.
26.6527.12	Alc-Subsc-Prof-Str	Subscriber profile string.
26.6527.13	Alc-SLA-Prof-Str	SLA profile string.
26.6527.16	Alc-ANCP-Str	ANCP string.
26.6527.17	Alc-Retail-Serv-Id	The service-id of the retailer to which this subscriber host belongs.
26.6527.18	Alc-Default-Router	The default gateway for the user (DHCP option [3] default-router for a DHCPv4 proxy)
26.6527.28	Alc-Int-Dest-Id-Str	Intermediate destination ID string.
26.6527.29	Alc-Primary-Nbns	The IPv4 address of the primary NetBios Name Server (NBNS).
26.6527.30	Alc-Secondary-Nbns	The IPv4 address of the secondary NetBios Name Server (NBNS).
26.6527.31	Alc-MSAP-Serv-Id	Service ID where the managed SAP is to be created.
26.6527.32	Alc-MSAP-Policy	Managed SAP policy used to create the MSAP.

Table 193 **AA-Answer Message — Accepted Authorization AVPs (Continued)**

AVP ID	AVP Name	Description
26.6527.33	Alc-MSAP-Interface	Group-interface name where the managed SAP is to be created.
26.6527.45	Alc-App-Prof-Str	Application profile string.
26.6527.99	Alc-Ipv6-Address	DHCPv6 IA-NA IPv6 address (wan-host).
26.6527.105	Alc-Ipv6-Primary-Dns	The IPv6 address of the primary DNSv6 server.
26.6527.106	Alc-Ipv6-Secondary-Dns	The IPv6 address of the secondary DNSv6 server.
26.6527.131	Alc-Delegated-Ipv6-Pool	The name of an IPv6 IA-PD prefix pool (pd-host).
26.6527.161	Alc-Delegated-Ipv6-Prefix-Length	DHCPv6 IA-PD prefix length (pd-host).
26.6527.174	Alc-Lease-Time	The lease-time for proxy, in seconds.
26.6527.181	Alc-SLAAC-IPv6-Pool	The name of an IPv6 SLAAC prefix pool (wan-host).
26.6527.1036	Alc-SPI-Sharing	grouped AVP Sets the SLA Profile Instance (SPI) sharing method for this subscriber session to SPI sharing per group or default.
26.6527.1037	Alc-SPI-Sharing-Type	Must be included in an Alc-SPI-Sharing grouped AVP. Sets the SPI sharing method. value 0 = default as specified in the SLA profile with **def-instance-sharing**. The Alc-SPI-Sharing-Id AVP should not be present. value 2 = per group; the group identifier is specified with the Alc-SPI-Sharing-Id AV.
26.6527.1038	Alc-SPI-Sharing-Id	Must be included in an Alc-SPI-Sharing grouped AVP. Specifies the group identifier when SPI sharing is per group.

16.6.1 Sample Configuration Steps

To specify the peers to reach the Diameter NASREQ server in a diameter peer policy:

```
configure
    aaa
        diameter-peer-policy "diameter-peer-policy-1" create
            description "Diameter NASREQ peer policy"
```

```
            applications nasreq
            origin-host "bng@nokia.com"
            origin-realm "nokia.com"
            peer "peer-1" create
                address 172.16.3.1
                destination-realm "myDSCRealm.com"
                no shutdown
            exit
        exit
```

To specify the Diameter NASREQ application specific parameters, such as AVP format and values, in a Diameter application policy:

```
configure
    subscriber-mgmt
        diameter-application-policy "diameter-nasreq-policy-1" create
            description "Diameter NASREQ application policy"
            application nasreq
            diameter-peer-policy "diameter-peer-policy-1"
            nasreq
                user-name-format mac
                include-avp
                    circuit-id
                    nas-port-id
                    nas-port-type
                    remote-id
                exit
            exit
        exit
```

To apply the Diameter NASREQ application policy as Diameter authentication policy at a VPLS capture SAP, at an IES/VPRN group-interface and/or at a local user database:

Note: A Diameter authentication policy cannot be configured simultaneously with a RADIUS authentication policy on the same group-interface or capture SAP, nor for the same host in a local user database.

```
configure
    service
        vpls 10 customer 1 create
            sap 1/1/4:*.* capture-sap create
                ---snip---
                diameter-auth-policy "diameter-nasreq-policy-1"
        ies 1000 customer 1 create
            subscriber-interface "sub-int-1" create
                ---snip---
                group-interface "group-int-1-1" create
                ---snip---
                    diameter-auth-policy "diameter-nasreq-policy-1"
        vprn 2000 customer 1 create
            subscriber-interface "sub-int-1" create
```

```
                ---snip---
                group-interface "group-int-1-1" create
                ---snip---
                    diameter-auth-policy "diameter-nasreq-policy-1"

configure
    subscriber-mgmt
        local-user-db "ludb-1" create
            ipoe
                host "ipoe-host-1" create
                    ---snip---
                    diameter-auth-policy "diameter-nasreq-policy-1"
            ppp
                host "ppp-host-1" create
                    diameter-auth-policy "diameter-nasreq-policy-1"
```

If no AA-Answer message is received from the primary or secondary Diameter peer, then the host or session can be instantiated with the configured defaults. This is achieved by the following NASREQ application policy configuration:

```
configure
    subscriber-mgmt
        diameter-application-policy "diameter-nasreq-policy-1" create
            on-failure failover enabled handling continue
```

To enable flexible integration with different NASREQ servers, a Python policy can be configured on the Diameter peer policy. The Python script can interact on the AVPs present in the AA-Request and AA-Answer messages.

```
configure
    python
        python-policy "py-policy-nasreq-1" create
            diameter aar direction egress script "NasreqAar"
            diameter aaa direction ingress script "NasreqAaa"
configure
    aaa
        diameter-peer-policy "diameter-peer-policy-1" create
            ---snip---
            python-policy "py-policy-nasreq-1"
```

16.7 Diameter Redundancy

Diameter redundancy is supported on multiple levels:

- High Availability (HA) refers to control plane redundancy with dual Control Plane Modules (CPMs) in a single chassis configuration. User sessions for Diameter applications are fully synchronized between CPMs.

  However, the TCP connections with peers are not synchronized and after a CPM switchover, need to be reestablished. Once the TCP peering connections are re-established, communication on the user level (NASREQ, Gx, or Gy) is resumed.
- Peer Level Redundancy — An existing user session can be rerouted to an alternate peer if the currently used peer fails.
- Multi-Chassis Redundancy — Diameter sessions are synchronized on the application level (through ESM if Gx and NASREQ) between two redundant SR OS nodes. Only one of those two nodes maintains a peering connection with the DRA/PCRF. As a result, the DRA or PCRF interacts with the redundant pair through a single peering connection.

16.7.1 Diameter Peer Level Redundancy

Each configured diameter node can support several peers that are simultaneously open. Only one of those peers is used to forward application messages for a given user session. If there are multiple peers for the same realm, then the next hop, peer is selected in the order of configured preference. If all peers are of the same preference for a given realm, then the peer index is used to break the tie.

Peer failover is performed by Diameter base protocol and is supported only for routable diameter request messages of the same user session (for example, peer failover does not apply to a CER message). A peer failover will trigger clearing of the destination host name from the user session, which will result in a Diameter routing lookup based on the destination-realm (since destination host name is cleared). A new next hop in the chain of priorities for a given destination realm within the realm routing table is selected.

Routable answer messages do not rely on peer failover procedures since their forwarding is governed on a hop-by-hop basis (exact reverse path of the request message).

A failover peer procedure for a user session is triggered under the following conditions related to the peer connectivity problems:

- Peering connection teardown — For example, a Diameter peer initiates a DPR message or sends a TCP FIN or RST message. In either case, the peer connection is closed, and a peer failover procedure is initiated.
- DWR timeout — Watchdog timeouts.

Once the failed peer is restored, forwarding is resumed on the restored peer for all new and existing sessions if this peer is the best next hop for the user session (dictated by peering and realm table). There is no user session stickiness to the peer and instead the routing of the Diameter request messages always follows the routing and forwarding tables. It is important to note that peer failover refers to the Diameter base ability to select a new best peer in the list of available peers, in the case where the current peer becomes unavailable. This rerouting of Diameter application messages is independent of the application level message retransmissions. Once the Diameter base determines that it cannot deliver the request message (for example, CCR times out or the next hop to the destination-realm is unavailable), the Diameter base notifies the application level. The application (NASREQ, Gx, or Gy) then decides (by a configuration option) whether it will re-transmit timed-out request messages (CCR).

16.7.2 Diameter Multi-Chassis Redundancy

The Diameter Multi-Chassis Redundancy solution on the node is based on the model where communication with the PCRF/DRA occurs over a single Diameter peering connection for the redundant pair of nodes; that is, an active Diameter proxy module running on the node is front-ending the communication with the PCRF/DRA, and is relaying messages between the Diameter clients (in redundant nodes) and the DRA/PCRF. Both nodes run the Diameter proxy module, but only the active one opens the TCP peering connection towards the DRA/PCRF.

The benefits of the Diameter proxy model are:

- The redundant pair of nodes appear as a single node to the PCRF/DRA.
- Diameter Identity is uniform across the pair of redundant nodes; that is, both nodes share the same origin-host/realm, a property that carries operational benefits.
- There is a robust and predictable failover of a single connection, as opposed to having separate diameter peering sessions per node.

The goal of this redundancy model is to provide a predictable and quick recovery after a node failure, PCRF failure, or relevant components within those two entities (such as line cards, MDAs, and physical ports).

Figure 239 illustrates the basic concept for Diameter Multi-Chassis Redundancy for a Gx application. The model shows two nodes (BNGs). Each BNG contains an ESM module and a Gx/Diameter module which have a peering connection to the active Diameter proxy module. The peering connections are IP connections. Both nodes communicate with the PCRF/DRA through the active Diameter proxy which maintains a peering connection with the PCRF/DRA.

Figure 239 **Proxy Gx Model**

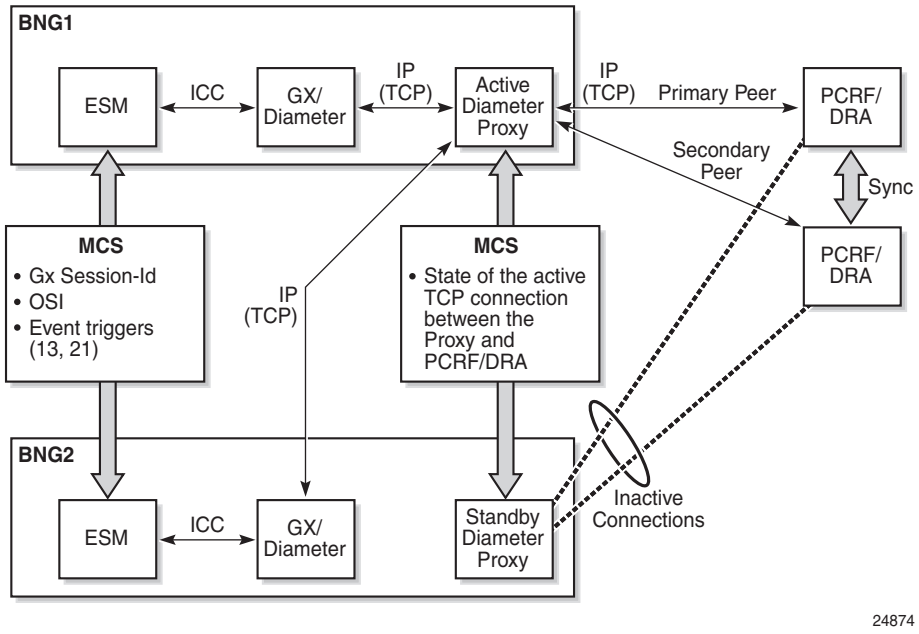

16.7.2.1 Diameter Proxy Model General Operational Principles

The fundamental principles of the Diameter proxy-based redundant solution are described below (also refer to Figure 239):

1. There can be only one active Diameter proxy per redundant pair of nodes. This active Diameter proxy maintains peering connections (primary/secondary) towards the PCRF/DRA at any given time. There are no peering connections open on the standby Diameter proxy. The active Diameter proxy accepts the connections from the client side, while the standby does not. The standby Diameter proxy ignores, without any reply, requests from the clients to open peering connections.

2. The activity selection is based on the system MAC address of each node and the current state of the Diameter proxy (Init, Active, Standby, Active-Wait, Standby-Wait and the Proxy-Switchover-Reg) in each node. This information is exchanged between the nodes via MCS. The system MAC address is unique per node. The activity selection is automatic and cannot be influenced via CLI, other than shutting down peers on the currently active Diameter proxy.

   Application level information (Gx session, NASREQ session, etc.) is not synchronized between Diameter proxies. Instead, synchronization of session level information is performed on the application level; for example, Gx session information is performed through subscriber management (ESM) synchronization.

3. Preemption is not supported if a node coming up from a bootup sequence would cause the MCS peer to transition from active to standby state. Only in the case where the two nodes are booting simultaneously, or are recovering from the MCS synchronization loss (both nodes are joining MCS while in the active state), would the node with a higher system MAC address transition into the active state (during the simultaneous node bootup) or remain in the active state (after isolation re-synchronization).
4. When the Diameter proxy transitions from an active to standby state, the newly-transitioned standby Diameter proxy closes the TCP connections towards the Diameter clients. The clients then initiate a new connection towards the secondary peer, but only after the current connection-timer expires (30 s by default). For example, assume that the connection timer is set for 30 s. If the TCP RST is received when the running connection-timer is at 15th s, then the secondary peer is not initiated for another 15 s.
5. In the case where the active Diameter proxy receives messages from the client while it does not have any peers on the server side open, the Diameter proxy sends a DIAMETER_UNABLE_TO_DELIVER (3002) message to the client. The client then retransmits the message to the same peer since it has only one peer that is in the UP state. Upon receipt of the DIAMETER_UNABLE_TO_DELIVER (3002) message, the client leaves the original message to time out, and then retransmits it (that is, the message is not retransmitted from the client side immediately upon the arrival of the DIAMETER_UNABLE_TO_DELIVER (3002) message).
6. The TCP protocol handles retransmissions on the transport layer to the same peer. This is valid on the Diameter client side and on the Diameter proxy side. A TCP retransmission normally occurs at the intervals driven by an exponential back-off. The initial timeout depends on the implementation, but for the most common case, it can be assumed to be 1.5 s, followed by an exponential back-off capped at 64 s (1.5 s, 3 s, 6 s, 12 s, 24 s, 48 s, 64 s).

7. The Diameter client handles retransmissions on the Diameter level. The Diameter client is able to retransmit on the same socket since it only has a single socket (to the active Diameter proxy). The T-bit in the Diameter header is set for every retransmitted message. The watchdog interval should be set to 1 s, so that the dead TCP connection (dead proxy) can be quickly identified.

   The Diameter client will retransmit when:

   - The original request times out.
   - It receives a reply with the E-bit set. Such retransmissions will not be triggered immediately upon the arrival of the response with E-bit set. Instead, the original messages that need to be retransmitted are left in the pending queue to time out and are retransmitted after the timeout period, which is controlled by the tx-timer command.
   - The primary peer is closed and the secondary peer is available.

8. The Diameter proxy never retransmits a message on the Diameter level since it does not perform any buffering (Tx/Pending queue). However, it does retransmit on the TCP level (hop-by-hop).
9. The Diameter proxy only relays messages between the client (application) and the server side (DRA/PCRF). The two bits in the Diameter header that the proxy is reacting on are the T-bit and the E-bit.

   If the T-bit is set in the message coming from the client side, the Diameter proxy sends the message to the secondary peer (invokes the peer failover procedure). That is, the application level retransmissions is performed by the Diameter client (which is peering with the Diameter proxy). The client sets the T-bit (retransmission bit) in the Diameter header and this signals to the Diameter proxy that it needs to failover the message to the alternate peer. This operation is performed on a per message basis and not on a per session basis.

   The Diameter proxy initiates a failover procedure to the secondary peer when the primary peer on Diameter proxy is closed, or the watchdog timer on the primary peer expires.

   All messages in the DRA/PCRF-to client direction with the E-bit set (the E-bit can be present only in answer messages) are dropped in the proxy. Consequently, the client retransmits the request, upon timeout.

   The messages with the E-bit set that are traveling in the opposite direction are not dropped; they are transparently passed to the DRA/PCRF.

10. On an SRRP Switchover, the AN-GW-IP of the newly-transitioned Master can be reported in CCR-u as an indication of the switchover. This functionality is enabled by arming the node with the event-trigger id 13 (USER_LOCATION_CHANGE) from PCRF.
11. On a Diameter proxy switchover, a SNMP Log/Trap is generated.
12. The standby client (SRRP standby) discards all messages received by the Diameter proxy.

16.7.2.2 Diameter Proxy Activity Selection

The Diameter proxy with the highest system MAC address assumes the controller role. The controller node decides which proxy becomes ACTIVE or STANDBY. Activity election information is processed by the controller node and then the controller node delegates the actual ACTIVE/STANDBY roles to Diameter proxies. The ACTIVE proxy may not necessarily be the same node as the controller node.

The activity selection (by the controller node) in the Diameter proxy is based on the current states of both Diameter Proxies (local and remote) and the system MAC.

Preemption is not supported, which means that newly brought up Diameter proxy does not overtake the activity state from the existing active Diameter proxy, regardless of the system MAC addresses.

Once the node becomes active, it advertises the new state to the MCS Diameter proxy peer and tries to open a DRA/PCRF peering connections and at the same time accept the client connections. The active Diameter proxy replies to the client with a DIAMETER_UNABLE_TO_DELIVER error-code when server side peers cannot be opened.

16.7.2.3 Synchronization and MCS

All application level (Gx or NASREQ) sessions related parameters are synchronized on the ESM level via MCS.

The parameters synchronized on the ESM level are:

- Session-Id
- Event-Triggers
- CC-Request-Number

The Diameter proxy module is synchronized via MCS; the information passed between the two nodes is:

- System MAC address: this address plays a role in the Diameter proxy state selection
- Controller MAC address: the System MAC address of the node that is performing the Diameter proxy state selection. The node with the highest system MAC address assumes the controller role. Once the controller makes the state selection for both nodes, it delegates those states to Diameter proxies. The controller role is collocated with one of the Diameter proxy nodes.
- Origin-State-Id (OSI)

- Diameter proxy States

The above information is used to determine the activity of the Diameter proxy at each node.

If an MCS link fails, the nodes become isolated. Each node acts independently and tries to become active. This scenario is described in Isolated Chassis.

16.7.2.4 Retransmissions

The handling of Diameter retransmissions is crucial for the Diameter Multi-Chassis Redundancy operation. Retransmissions provide the means to recover a Diameter session that was left in an unacknowledged state due to failure of the path between the node and the DRA/PCRF.

Retransmissions of Diameter messages are handled on two levels by a pair of redundant nodes:

1. At the TCP level: request and answer messages are retransmitted by TCP. These types of retransmissions are only significant between two directly connected peers, hop-by-hop retransmissions. For example, if the failure occurs beyond two directly connected peers, these type of retransmissions will not help.
2. At the application (Gx, Gy, NASREQ) level: only request messages are retransmitted. These types of retransmissions extend beyond two directly connected peers and can cover end-to-end failure cases.

A more detailed explanation of the processing that occurs on each level for a Gx application is given below.

1. The first level of retransmissions occurs at the TCP level. The Gx message is handed over to the TCP socket Base Diameter module. TCP tries to deliver this message in a connection-oriented (reliable) fashion. If a TCP ACK for the transmitted message is not received, the message, in the most common case, is retransmitted in intervals of 1.5 s, 3 s, 6 s, up to 64 s. After 10 s of trying to retransmit the message to the same peer at the TCP level, the Base Diameter tries to retransmit (configuration dependent) the message to the next TCP socket (secondary peer). The assumption is that the primary peer is unavailable (busy, failed, or the network path to it is broken) after 10 s of trying. The TCP retransmissions are peer oriented and very localized (to the particular TCP connection on the particular BNG node). If a network fails, TCP retransmissions cannot re-route the traffic to an alternate destination. As such, they cannot

protect against the peer (PCRF) failure or the BNG node failure. They can, however, indirectly provide a clue that something is happening at the peer level, so that the upper layers can take adequate actions. Watchdogs are also used to detect peer failure and they can provide faster detection of the peer failure (after 2 s).

2. Re-routing of the traffic occurs at the Application/Diameter level. Once the peer is considered unavailable, or the original requests message times out, the Diameter has the ability to re-route the retransmitted message to an alternate (secondary) peer. The Diameter level retransmissions can protect against a PCRF/DRA failure. Traffic can be switched to the secondary peer (this functionality must be enabled via configuration).

   Since the Diameter proxy only relays the messages between the client and the DRA/PCRF, it never buffers and retransmits the Diameter message. Retransmissions are the responsibility of the Diameter peer (Diameter client) that sits behind the Diameter proxy. Retransmitting Diameter client sets the T-bit in the Diameter header of the retransmitted message (CC-req-num is kept the same in the original and retransmitted message). The T-bit in the message triggers the Diameter proxy to re-route the messages to the secondary peer while the primary peer is still active. This means that the Gx client has already retransmitted the message, and the Diameter proxy re-route its.

   With single peer, the Diameter **client** retransmits the message to the same peer and it sets the T-bit in the Diameter header.

   With a single peer, the Diameter **proxy** sends the message with the T-bit set to the same peer.

   The Gx client typically re-routes the message to the secondary peer in the following cases:

   – The primary peer is closed and the secondary peer is available.

   – The Diameter message times out due to no response.

   – The Diameter message times out due to a received answer with E-bit set in the Diameter header (E-bit can be only set in the answer messages and it indicates protocol errors with Result-Code from the 3xxx range). Once the reply with E-bit set is received, the corresponding request message is left on the pending queue where it times out after the interval controlled by the tx-timer statement in the diameter-application-policy. Upon message timeout, the node retransmits the message to the secondary peer if the secondary peer is available. If not, the message will is retransmitted to the same (primary) peer.

In summary:

- TCP handles retransmissions towards the peer (hop-by-hop retransmissions).

- Diameter and Diameter applications (Gx, Gy, NASREQ) retransmit to the secondary peer in cases where the application level message times out, a protocol error is received (Result-Code 3xxx) in the answer from the DRA/PCRF, or the primary peer is 'closed'. If there is only one peer available (primary), the Diameter application retransmits to that peer. The T-bit in retransmitted messages is always set. Diameter level retransmissions cover failure cases that extend beyond two directly connected hops.
- The Diameter proxy never retransmits (retransmissions are handled by the Diameter client that sits behind the proxy). However, the Diameter proxy sends messages with the T-bit set to the secondary peer.

These scenarios are shown in Figure 240, Figure 241, and Figure 242.

Figure 240 **Retransmissions with Two Peers and No Diameter Proxy**

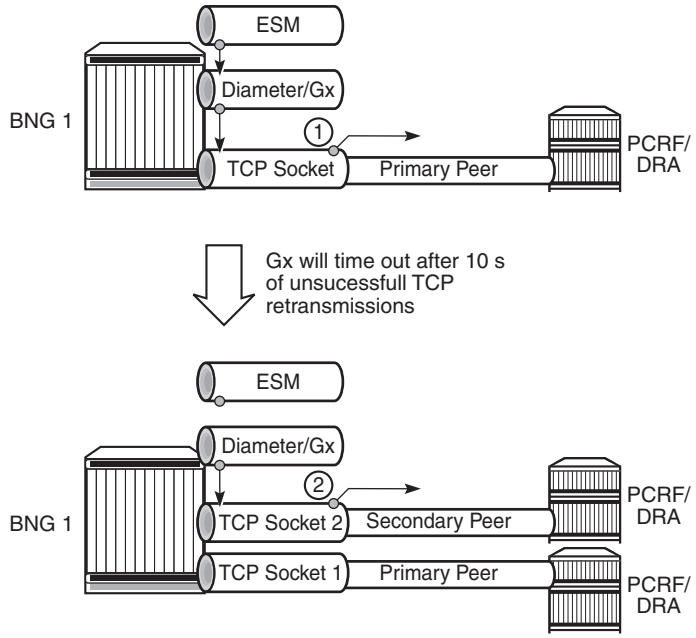

(1) Req/Answ retransmissions governed by TCP (such as windowing and exp backoff).

(2) Request retransmissions handled at the Gx level (peer failover). T-bit is set.

24879

Figure 241 **Retransmissions with a Single Peer and No Diameter Proxy**

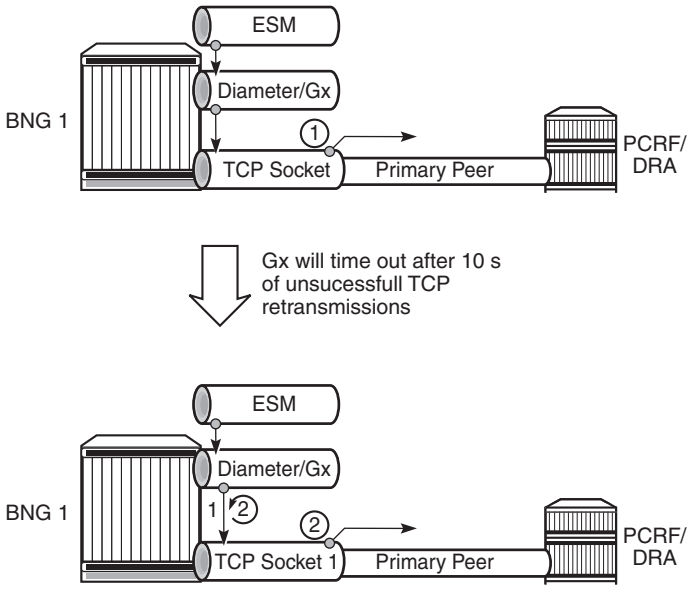

(1) Req/Answ retransmissions governed by TCP (such as windowing and exp backoff).

(2) Request retransmissions handled at the Gx level (peer failover). T-bit is set.

24880

Figure 242 **Redirection with Diameter Proxy (T-bit set)**

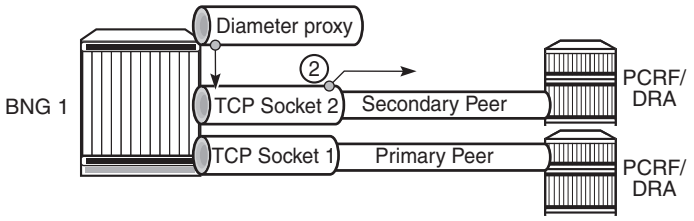

(2) Diameter Proxy will use secondary peer only if T-bit is set, or the primary peer is closed.

24881

16.7.2.5 Retransmissions and the T-bit

Multi-chassis redundancy is less concerned with the retransmissions of the answer messages (RAA) since, if the answer is not received, the PCRF retransmits the request (RAR). Retransmission of the **answer** messages is performed only on the TCP level within a single
node. It is not performed on the Diameter level.

When the **request** message is retransmitted by the Diameter application (due to the Tx timer timeout, primary peer failure - DWR timeout, or receipt of the answer message with E-bit set), the content of the message stays the same, including the CC-Request numbers but the T-bit in the Diameter header is set. The T-bit indicates to the PCRF that the message is retransmitted (mostly used for accounting purposes so that the counting records are not duplicated). It also signals to the Diameter proxy that the message rerouting to the secondary peer should be performed.

16.7.2.6 Diameter Proxy Role

The Diameter proxy is applied to an IPv4 or IPv6 address within a routing-context on a node. This IP address is a Diameter proxy listening IP address that is associated with an interface on a node, including the system interface (system IPv4 or IPv6 address) or loopback interface (loopback IPv4 or IPv6 address).

The number of Diameter Proxies per listening IP address is limited to one. That is, each proxy diameter-peer-policy requires a unique combination of source-ip (listening IP address) and the routing-context (node).

The number of Diameter-peer-policies on a node is limited to 32. This means that the combined number of Diameter clients and Proxies on a node cannot exceed 32.

The Diameter Proxy has the following role on a node:

1. The active Diameter proxy relays messages between the application (Gx and NASREQ) module on a node and the PCRF/DRA.
2. Only the active Diameter proxy allows peering connection with the client (the Diameter on a node). The standby Diameter proxy refuses the client connection.
3. The Diameter proxy with open peering connections is referred to as the active Diameter proxy (ADP). Its counterpart is called the standby Diameter proxy.
4. The Diameter proxy retransmits the message on the TCP level towards the same peer.

5. However, the Diameter proxy does not perform application level message retransmission and the peer failover procedure due to the timeout of the application level message. Instead, the application level retransmission is performed by the Diameter client (which is peering with the Diameter proxy). The client sets the T-bit (retransmission bit) in the Diameter header of the retransmitted message (same CC-Req_Number) and this signals to the Diameter proxy that it needs to failover the message to the alternate peer.

Note: This re-routing operation in the proxy is performed per message and not per sessions, as it is the case for a Diameter client.

6. When the primary peer is closed, or the watchdog timer on the primary peer expires, the Diameter proxy initiates failover procedure to the secondary peer.
7. The standby client (SRRP standby) discards received RAR messages.
8. The active Diameter proxy replies to the client with the error message UNABLE_TO_DELIVER if the peering connection towards the server cannot be open. The client retransmits the messages (since it has only one connection) after the timeout interval.

   Table 194 summarizes the differences between the regular Diameter client and the Diameter proxy.

Table 194 **Summary of Differences**

Regular Diameter Client	Diameter Proxy
Initiates messages.	Transparently passes all messages between the client and the server. Never initiates the messages.
Buffering is implemented, thus retransmissions are supported.	Buffering is not implemented (pending queue), thus messages are never retransmitted.
When retransmitting, it sets the T-bit in the Diameter Header.	Never retransmits the messages.
Failover to the secondary peer is triggered by: • Message timeout • Primary peer is down; immediate failover. • All messages with an E-bit set triggers failover after the message times out on the pending queue.	Failover to the secondary peer is triggered by: • Messages with T-bit set; immediate failover per message. • Primary peer shutdown
Diameter client performs peer failover per session.	Diameter proxy performs peer failover per message (with the T-bit set).

16.7.2.7 Diameter Proxy and CC-Request-Number AVP

CC-Request-Number AVP (RFC 4006, 8.2) are typically used to match requests with answers. Session-id and CC-Req-Num are a unique per-message pair. CC Request Numbers along with the session-id uniquely identify a transaction (matching requests and answer messages) on a global level.

The Diameter proxy does not re-write the CC-Request-Number in the messages received by the client.

CC-Req-numbers are synchronized at the ESM level. This is needed so that operation with proper CC-Req-Num can resume after the switchover.

For example, the following CC-Req-Num sequence for the session is preserved across SRRP switch-overs:

- 1st host (e.g. IPv4) of a dual-stack host is setup
- CCR-I with CC-Req-Num = 1 is sent
- 2nd host (for example,. IPv6 IA-NA) of the same dual-stack host is set up
- CCR-U with CC-Req-Num=2 is sent
- SRRP switchover occurs
- 2nd host in the dual-host is removed (DHPCv6 release)
- CCR-U with CC-Req-Num=3 is sent from the new SRRP Master

16.7.2.8 Stateless Diameter Proxy

The Diameter proxy does not maintain any session state. Forwarding is based on transactions which are short lived. Transactions are based on a pairing request/ answer messages matched by the same hop-by-hop identifier and the peer from which the request was received. In this fashion, answer messages coming from the DRA/PCRF can be unambiguously forwarded to the proper Diameter client (from which the request was received).

Since the session state is not kept in Diameter proxy, RAR request are be flooded to both Diameter clients.

16.7.2.9 Switchover Scenarios

The following are four types of switchovers that are most likely to occur:

1. Switchover to a new DRA/PCRF peer at a Diameter proxy
2. Diameter proxy switchover due to failed peers on the server side
3. Diameter proxy switchover due to Diameter proxy node failure
4. SRRP switchover

Each switchover type for a Gx application is discussed in more detail below:

1. Switchover to a new PCRF/DRA peer is handled at the Diameter proxy level. This scenario is shown in Figure 243.

   The Diameter proxy switches over to the new peer in the event of two cases:

   – It receives a Diameter message with the T-bit set (retransmission bit in the Diameter header) from the Diameter client.

   Retransmission due to the message timeout is performed at the Diameter **client** level, and setting the T-bit signals to the Diameter proxy that peer failover is needed for this particular message.

   – The TCP connection to the primary peer is explicitly closed (for example, due to
   TCP RST or watchdog timeouts). Then, the Diameter proxy performs a fail-over of all sessions to the secondary peer immediately.

Figure 243 **PCRF/DRA Peer Switchover**

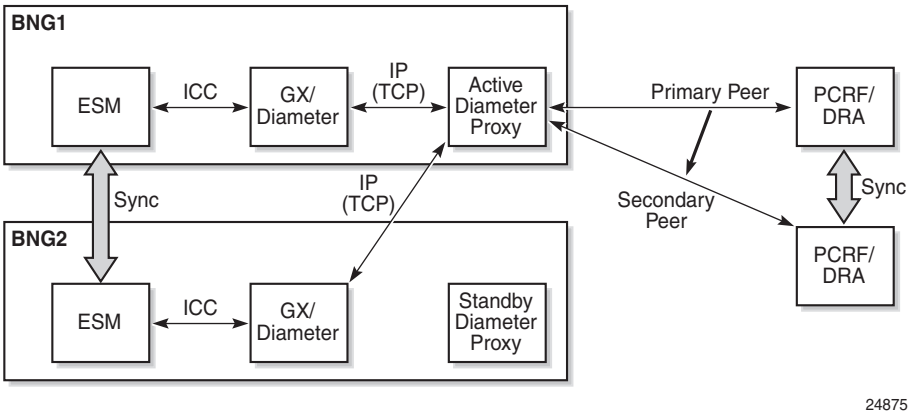

2. Activity switchover between the Diameter proxies occurs when the active Diameter proxy loses all peering connections to the PCRFs/DRAs while both nodes (including MCS) are operational.

This type of switchover (without the node failure) is unlikely to occur. For example, a Diameter proxy switchover (without the node failure) would mean that all PCRFs/DRAs have failed, since normally the same pair of PCRF/DRA peers are configured in both Diameter proxy nodes. This would mean that all DRA/PCRFs are unavailable, which indicates a problem on the network side (either network paths to DRA/PCRFs are broken or all DRA/PCRFs have crashed).

Another example where this scenario could occur is poor redundancy design practices. For example, the active Diameter proxy has a single peering connection to one PCRF (no secondary peer), while the standby Diameter proxy is configured with a separate PCRF peer (the two PCRFs are still synch'd). This scenario is shown in Figure 244.

Figure 244 **Diameter Proxy Switchover**

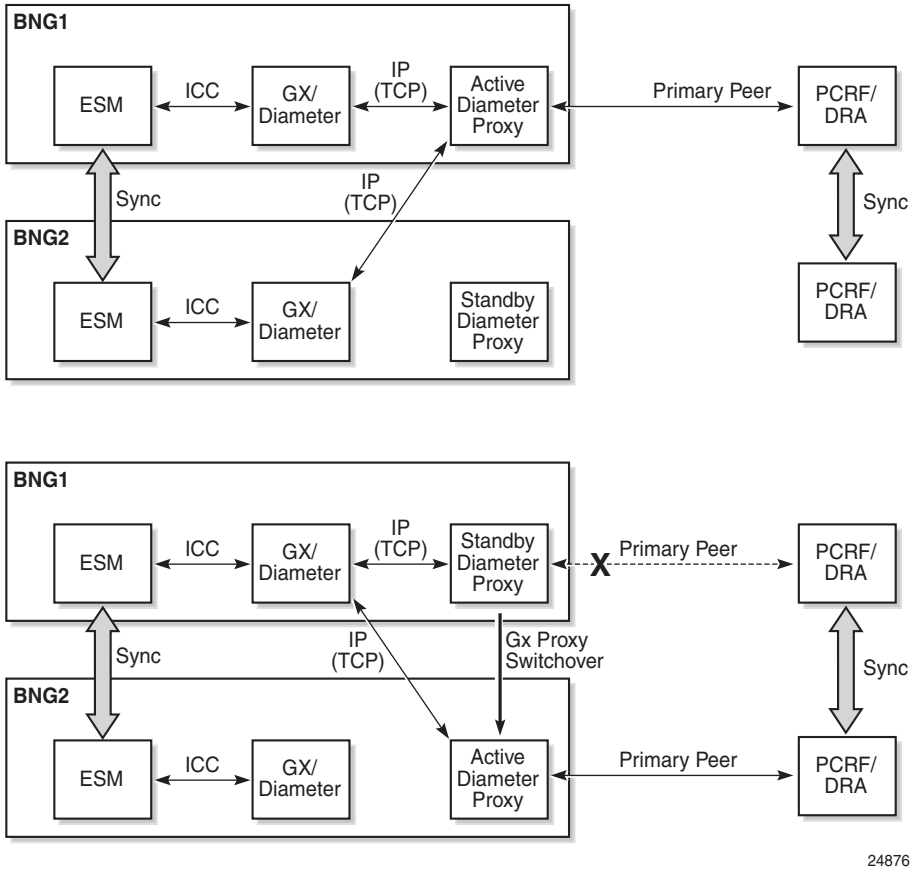

As always, only the active Diameter proxy maintains an open TCP peering connection towards the PCRF/DRA. If this connection fails, the active proxy sends a TCP RST towards the client and transitions into standby state. The client then, upon expiry of connection-timer, open a new TCP connection towards the newly active Diameter proxy.

3. In this scenario, the entire proxy node fails, as shown in Figure 245. The surviving node resumes operation.

Figure 245 **Node Failure**

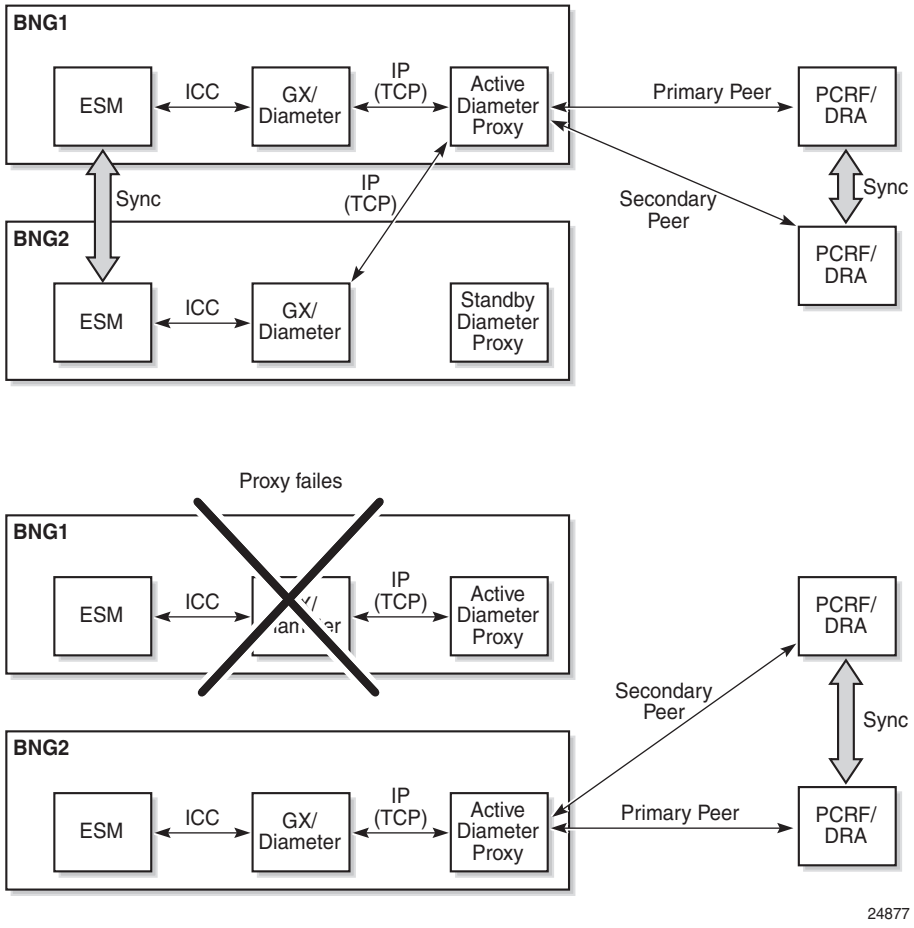

4. The last scenario is where the switchover occurs in the access and the SRRP switches activity. The new SRRP master resumes uninterrupted operation.

16.7.2.10 Log/Trap Generation Caused by Diameter Proxy State Change

In cases where the Diameter proxy changes its states (INIT, ACTIVE, STANDBY), a log/trap is generated. This log is enabled by default in log-event control. The notification name is tmnxDiamProxyStateChange.

16.7.2.11 Switchover Update Event (CCR-u)

The AN-GW-Address carried in the CCR-I message for the Diameter application session (for example, Gx) is the IP address of the node on which the underlying SRRP instance (for this Gx session) is in the SRRP master state.

When the SRRP switches over due to the failure in the access part of the network (including the ports on a node), a CCR-U can be optionally (configuration dependent) sent with the AN-GW-Address AVP of the node on which an SRRP instance transitioned into the master state. This behavior is controlled via the event trigger id 13 (USER_LOCATION_CHANGE).

16.7.2.12 Isolated Chassis

In cases where the MCS connection is broken, the Diameter proxy on both nodes try to become active since they each consider that they are the only functional node. From the local point of view, the MCS peer is dead.

While in isolation scenario, both nodes are most likely able to open the TCP peering session with the PCRF/DRA (see Figure 246).

Figure 246 **Isolated Nodes**

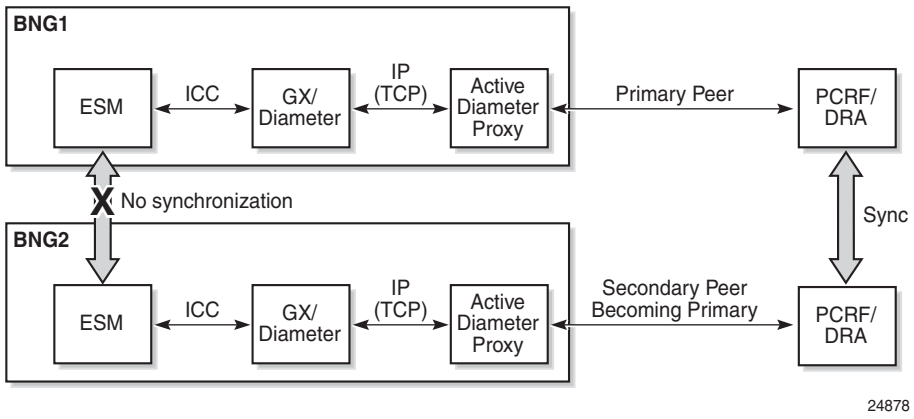

Once the MCS is recovered, the states are re-synchronized.

16.7.2.13 Diameter Identities

Diameter identities (origin-host/realm) can be configured to be the same on both nodes. This ensures that the redundant pair of nodes appears as a single node at the Diameter level (Diameter Identities).

16.7.2.14 High Availability

A CPM switchover on the active Diameter proxy causes the peering connections between the client and the proxy to be lost. Consequently, the clients have to re-establish their peering connections. Peering connections on the active Diameter proxy towards the server remain uninterrupted.

16.7.3 Gx Specific Behavior

Gx specific behavior in a multi-chassis configuration is as follows:

1. A Gx client at a node attempts to open a TCP connections to both Diameter proxies, but only the active Diameter proxy accepts accept the TCP request (see Diameter Multi-Chassis Redundancy).

2. The standby Diameter proxy ignores the connection request and does not respond in any way (not even TCP RST).
3. The Gx client normally tries to reopen the configured connections (peers) every connection-timer interval (30 s by default).
4. Since the Gx client has only one peering connection open, retransmissions due to the application level message timeout occurs on that same peer. The T-bit is set and signals to the Diameter proxy that it needs to perform a peer failover procedure.
5. The Gx client discards/ignores all messages received on the standby node (standby SRRP).

16.8 Diameter Debugging

This section explains the debugging capabilities for troubleshooting a Diameter or Diameter application problem.

Debugging should be as specific as possible and limited to relevant messages only. Enabling detailed debugging for all diameter messages on a production node generates a flood of information and is not very helpful in isolating the problem; it can, however, be a valid scenario for lab testing.

Diameter debug output can be limited to:

- messages to a specified origin and/or destination realm

  Messages that have no origin or destination realm AVP to match against will also be shown in this configuration.

```
debug
    diameter
        dest-realm dest.com
        origin-realm origin.com
    exit
exit
```

- messages of specified message types

```
debug
    diameter
        message-type ccr cca
    exit
exit
```

- messages sent or received on peers belonging to a specific Diameter peer policy. Up to eight diameter peer policies can be specified

```
debug
    diameter
        diameter-peer-policy "diameter-peer-policy-1"
        exit
        diameter-peer-policy "diameter-peer-policy-2"
        exit
    exit
exit
```

- messages sent or received on a specific Diameter peer

```
debug
    diameter
        diameter-peer "peer-1"
    exit
exit
```

Only messages that match all specified criteria are shown in the debug output.

The **diameter debug detail-level** command can be set to

- Low (default): displays the Diameter message header
- Medium: displays the Diameter message header and AVPs
- High: displays the Diameter message header, AVPs, and a packet hex dump

```
debug
    diameter
        detail-level medium
    exit
exit
```

For a per-diameter peer policy additional criteria can be specified to further restrict the debug output to:

- messages sent or received on a specific Diameter peer

  overrides the peer configuration at the debug diameter level

```
debug
    diameter
        diameter-peer-policy "diameter-peer-policy-1"
            diameter-peer "peer-1"
        exit
    exit
exit
```

- messages of specified message types

  overrides the message type at the debug diameter level

```
debug
    diameter
        diameter-peer-policy "diameter-peer-policy-1"
            message-type ccr cca rar raa
        exit
    exit
exit
```

- messages from a diameter session where the diameter session is learned from Diameter application messages matching specified AVP values (avp-match criteria)

```
debug
    diameter
        diameter-peer-policy "diameter-peer-policy-1"
            avp-match 1
                avp 443.444 octetstring ascii value "1/1/1:10.20@bng1"
                message-type ccr
                no shutdown
            exit
        exit
    exit
exit
```

Only messages sent or received on a peer that belongs to the diameter peer policy and matches all specified criteria are shown in the debug output.

To restrict debug output to messages from a given diameter session, AVP matching with session learning is required. AVP 263 Session-Id is the only Diameter AVP that is present in all messages of a Diameter session and is dynamically assigned when the Diameter session is initiated. The session ID is not useful to specify as criteria in debugging because its value is unknown upfront. Instead, using AVP matching, it is possible to learn the session ID from Diameter application messages (NASREQ, Gx, and Gy) matching a message type and one or multiple AVP values. All subsequent Diameter application messages that belong to the learned session ID are then included in the debug output, as shown in the following example:

```
debug
    diameter
        detail-level medium
        diameter-peer-policy "diam-pol-1"
            avp-match 1
                avp 25 octetstring ascii value "trace-me"
                avp 10415-1005 octetstring ascii value "sub-id:vip-1"
                message-type cca
                no shutdown
            exit
            avp-match 2
                avp 443.444 octetstring ascii value "1/1/1:10.20@bng1"
                message-type ccr
                no shutdown
            exit
            message-type ccr cca rar raa
        exit
        message-type all
    exit
exit
```

In this configuration, the debug output for Diameter messages is restricted to messages of type ccr, cca, rar, or raa that are sent or received on diameter peer policy “diam-pol-1” and that belong to Diameter sessions for which:

- a CCA message is received with Class AVP = “trace-me” and 3GPP Charging-Rule-Name AVP = “sub-id:vip-1”. The session ID is learned from the message. All subsequent ccr, cca, rar, and raa messages with the same session ID are also shown until a relearning occurs: a new CCA message received matching the AVP criteria specified for this AVP match ID. If still available in the system, the corresponding CCR message is also shown.
- a CCR message is received with subscription-id-data AVP = "1/1/1:10.20@bng1" (grouped within the subscription-id AVP). Session ID learning occurs for this AVP match ID and all subsequent ccr, cca, rar, and raa messages with the same session ID are shown as described above.

The rules for Diameter debug AVP matching are summarized below:

- Diameter debug AVP match criteria are specified in an **avp-match** *id* command. Up to five AVP match IDs can be specified per diameter peer policy. For each AVP match ID, a unique Session ID is learned and the corresponding messages are logged in debug: OR function for all **debug diameter diameter-peer-policy** *policy* **avp-match** *id* commands.
- For each AVP match ID, a message-type and at least one AVP ID criteria must be specified. Up to five AVP ID criteria can be specified per each AVP match ID. Session ID learning or relearning occurs when all specified AVP ID criteria are matched in any of the specified application message types: AND function for all **debug diameter diameter-peer-policy** *policy* **avp-match** *id* **avp** *avp-id* commands
- For each AVP in an AVP match, the AVP ID, type, and value should be specified:
  - *avp-id* is specified as [*vendor-id*-]*avp-code*[.*avp-id*] with nesting up to five levels deep. For example: “443.444” to indicate Subscription-Id. Subscription-Id-Data and “456.10415-871” to indicate Multiple-Services-Credit-Control.Quota-Holding-Time.
  - *type* corresponds to the type as defined in the standards or as defined in the *Diameter Interface Specification* document for vendor-specific AVPs. The following types are supported: octetstring (used for utf8string, octetstring and diameterid), integer32, integer64, unsigned32, unsigned64, and address. Any AVP can be specified as an octetstring in hex format.
  - *value* is the AVP value to be matched. Optionally the format in which the value is specified can be configured as ascii or hex. A hex string is specified as "0xaabbcc…" (starting with "0x" followed by the hex nibbles).
- The message type specified in an **avp-match** *id* command is only used as match criteria and does not limit the debug output:

```
debug
    diameter
        diameter-peer-policy "diam-pol-1"
            avp-match 1
                avp 443.444 octetstring ascii value "1/1/1:10.20@bng1"
                message-type ccr
                no shutdown
            exit
            message-type ccr cca rar raa
        exit
    exit
exit
```

In this configuration, the session ID is learned only from CCR messages with subscription-id data AVP = “1/1/1:10.20@bng1”. When the session ID is learned, all CCR, CCA, RAR and RAA messages from that session ID are shown in the debug output.

- When at least one AVP match ID is configured with the **debug diameter diameter-peer-policy** *policy* **avp-match** *id* **no shutdown** command, only the messages for learned session IDs for the specified diameter peer policy are shown in the debug output. An AVP match ID in shutdown state is not active (session ID learning disabled and no debug messages are logged for this AVP match ID).
- If the session ID learning happens based on a match in an Answer message, then the corresponding Request message is also logged in debug only if it was still available in the system.
- Relearning occurs when a message from a session ID different from the learned session ID matches all specified AVP ID criteria in any of the specified application message types. The learned session ID for the AVP match ID is overwritten.
- Messages from a learned session ID are no longer matched against the AVP match criteria. The same session ID is learned only once.
- Learning or relearning events are reported in the debug output:

```
174 2016/01/21 15:31:02.63 UTC MINOR: DEBUG #2001 Base DIAMETER
"DIAMETER: Session-Id learning
diameter-peer-policy diameter-peer-policy-1 avp-match 1 learned Session-
Id bng.origin.com;1453294062;8 (replacing Session-Id bng.origin.com;1453294062;7)"
```

- The learned session ID is deleted when the corresponding AVP match ID debug command is deleted or shutdown.
- The **tools>dump>debug>diam>avp-match-learned-session-id** command displays the current learned Diameter session IDs for each Diameter application policy AVP match ID.

17 Python Script Support for ESM

17.1 Python Script Support for ESM

In order to provide programmable flexibility in ESM applications, the SR OS provides the following features with Python script support:

- sub-ident-policy
- radius-script-policy
- python-policy

17.2 Python in SR OS Overview

The SR-OS python script support is based on Python version 2.4.2. Python has a set of language features (such as functions, lists and dictionaries) and a very large set of packages which provide most of the Python functionality. By keeping the language features intact and drastically reducing the number of packages available, the operator is provided with a flexible, although small, scripting language.

The only feature removed from the Python language is unicode support. The only packages provided to the operator are:

- **alc** — The SR OS-provided packages provide access to various ESM objects such as DHCPv4, DHCPv6 or RADIUS packets
- **binascii** — Common ASCII decoding like base64
- **re** — Regular expression support
- **struct** — Parses and manipulates binary strings
- **md5** — MD5 message digest algorithm

17.2.1 Python Policy – GTPv1-C API

The system will provide a Python object for input GTPv1-C packet alc.gtp1c. The alc.gtp1c packet has following attributes to represent the GTPv1-C header fields as displayed in Table 195.

Table 195 **GTPv1-C Header Fields**

Class Attrs	GTPv1-C Header Field	Access
alc.gtp1c.version	Integer, version field in the header	Read
alc.gtp1c.ptbit	boolean, true mean 1, false means 0; PT bit in the header	Read
alc.gtp1c.ebit	boolean, E bit in the header	Read
alc.gtp1c.sbit	boolean, S bit in the header	Read
alc.gtp1c.pnbit	boolean, PN bit in the header	Read
alc.gtp1c.type	Integer message type field in the header	Read
alc.gtp1c.len	Integer message length field in the header	Read
alc.gtp1c.teid	Unsigned integer, or None if TEID does not exist	Read
alc.gtp1c.seq	Integer sequence number in the header	Read
alc.gtp1c.npdu_number	Integer N-PDU Number field in the header	Read
alc.gtp1c.next_ext_type	Integer Next Extension Type field in the header	Read
alc.gtp1c.ext_list	Tuple, a tuple of extension headers; each element in the tuple represent one extension header: (ext_type, ext_content)	Read

The following is a list of alc.gtp1c functions:

- **alc.gtp1c.drop()**: The system drops the resulting packet.
- **alc.gtp1c.getIEList()**: The system returns a tuple of ie-type, each element represent one IE instance in the packet.
- **alc.gtp1c.get(ie_type)**: The system returns a tuple of str, each str is the IE data of specified ie_type instances. If the specified IE doesn't exist, then return (); if the IE length=0, then return "" for this IE instance.
- **alc.gtp1c.set(ie_type, ie_tuple)**: This function removes all instances of ie_type and insert a list of new IE instances. Each element in ie_tuple is a string, represent one instance of IE to be inserted. For each inserted IE instance, the type is ie_type. System will use TV format for type<128 and TLV format for type>=128.
- **alc.gtp1c.clear(ie_type)**: This function removes all existing instances of ie_type.

17.2.2 Python Policy – GTPv2-C API

The system will provide a Python object for input GTPv2-C packet: alc.gtp2c.

Shown in Table 196, the alc.gtp2c packet has following attributes to represent the GTPv2-C header fields:

Table 196 **GTPv2-C API**

Class Attrs	GTPv2-C Header Field	Access
alc.gtp2c.version	Integer version field in the header	Read
alc.gtp2c.pbit	boolean, true mean 1, false means 0; P bit in the header	Read
alc.gtp2c.tbit	boolean, T bit in the header	Read
alc.gtp2c.type	Integer message type field in the header	Read
alc.gtp2c.len	Integer message length field in the header	Read
alc.gtp2c.teid	Unsigned integer, or None if TEID doesn't exist	Read
alc.gtp2c.seq	Integer sequence number in the header	Read

The following is a list of alc.gtp2c functions:

- **alc.gtp2c.drop()**: The system drops the resulting packet.
- **alc.gtp2c.getIEList()**: The system returns a tuple of 2-element tuple (ie-type, ie-instance). Each 2-element tuple represent one IE instance in the packet. For example: a packet with the following IEs:
  - type:2, instance: 1
  - type:2, instance: 1
  - type:74, instance: 0

  The system will return ((2,1),(2,1),(74,0))
- **alc.gtp2c.get(ie_type, instance_id)**: The system returns a tuple of str. Each str is the IE data of specified (ie-type, instance-id). If the IE is a grouped IE, then the str is also the content of whole grouped IE. If the specified (ie_type, instance_id) doesn't exist, then return (); if the IE length=0, then return "" for this IE instance.

- **alc.gtp2c.set(ie_type, instance_id, ie_tuple)**: This function removes all instances of (ie_type, instance_id) and inserts a list of new IE instances. Each element in an ie_tuple is a string, represent one instance of IE to be inserted. For each inserted IE instance, the type is ie_type, the instance is instance_id.
- **alc.gtp2c.clear(ie_type, instance_id)**: This function removes all existing instances of (ie_type, instance-id).
- **alc.gtp2c.str_to_groupedIE (parent_ie_str)**: Parses the parent_ie_str and return a tuple of embedded IEs. Parent_ie_str is a string of data field of parent IE; each element of returned tuple is a tuple of (ie_type, ie_instance, ie_data), which represent one embedded IE. For example, with the following code:
  - ie_str = alc.gtp2c.get(93,0)[0]
  - ie_tuple = alc.gtp2c.str_to_groupedie(ie_str)

  The above code parses the first instance of the bearer context with an instance=0. Assume the bearer context looks like that shown in Figure 247:

Figure 247 **Bearer Context Example**

```
⊞ Bearer Context : [Grouped IE]
    IE Type: Bearer Context (93)
    IE Length: 44
    0000 .... = CR flag: 0
    .... 0000 = Instance: 0
  ⊟ EPS Bearer ID (EBI) : 5
      IE Type: EPS Bearer ID (EBI) (73)
      IE Length: 1
      0000 .... = CR flag: 0
      .... 0000 = Instance: 0
      0000 .... = Spare bit(s): 0
      .... 0101 = EPS Bearer ID (EBI): 5
  ⊟ Fully Qualifed Tunnel Endpoint Identifier (F-TEID) : S2b-U ePDG GTP-U interface, TEID/GRE Key: 0xfe500225, IPv4 10.190.172.254
      IE Type: Fully Qualified Tunnel Endpoint Identifier (F-TEID) (87)
      IE Length: 9
      0000.... = CR flag: 0
      .... 0101 = Instance: 5
      1... .... = V4: IPv4 address present
      .0.. .... = V6: IPv6 address not present
      ..01 1111 = Interface Type: s2b-u ePDG GTP-U interface (31)
      TEID/GRE Key: 0xfe500225
      F-TEID IPv4: 10.190.172.254 (10.190.172.254)
  ⊟ Bearer Level Quality of Service (Bearer QoS) :
      IE Type: Bearer Level Quality of Service (Bearer QoS) (80)
      IE Length: 22
      0000.... = CR flag: 0
      .... 0000 = Instance: 0
      .... ...0 = PVI (Pre-emtion Vulnerability): Enabled
      ..00 01.. = PL (Priority Level): 1
      .0.. .... = PCI (Pre-emtion Capability): Enabled
      Label (QCI): 8
      Maximum Bit Rate for uplink: 0
      Maximum Bit Rate for Downlink: 0
      Guaranteed Bit Rate for uplink: 0
      Guaranteed Bit Rate for Downlink: 0
```

1043

Then, the ie_tuple is

((5,0,'\x05'),(87,5'\x9f\xfe\x50\x02\x25\x0a\xbe\xac\xfe'),(80,0,'\x04\x08\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00'))

An exception is raised if the system cannot parse the parent_ie_str into a ie_tuple.

- **alc.gtp2c.groupedie_to_str(grouped_ie_tuple)**: Converts the grouped_ie_tuple into a str and return the str (see the description of alc.gtp2c.str_to_groupedie for the format grouped_ie_tuple).
  An exception is raised if the conversion fails. The following is an example:

  ie_str = alc.gtp2c.get(93,0)[0]

```
ie_tuple = alc.gtp2c.str_to_groupedie(ie_str)
fteid = list(ie_tuple[1])
fteid[1]=0
ie_str = alc.gtp2c.groupdedie_to_str((ie_tuple[0],tuple(ftied),ie_tuple[2]))
alc.gtp2c.set(93,0,(ie_str))
```

With previous example packet, the code changes the instance of the embedded F-TEID to 0.

17.2.3 Python Changes

The following changes have been made to Python for it to run on an embedded system:

- No files or sockets can be opened from inside Python scripts.
- No system calls can be made from inside Python scripts nor is the posix package available.
- The maximum recursion depth is fixed to twenty.
- The total amount of dynamic memory available for Python itself and Python scripts (excluding the Python cache) is capped at 8 Mb for 32-bit SR OS and 16 Mb for 64-bit SR OS on a CPM.
- The size of the Python script source file should be less than 60 kb.

17.3 Python Support in sub-ident-policy

A Python script can be configured in **sub-ident-policy** to return following ESM attributes:

- sub-id
- sla-profile name
- sub-profile name
- ancp-string

The system runs the Python script configured in the sub-ident-policy against the received DHCPv4 ACK message. This is used as the input of the script. Within the script, the user can set the value with the above ESM attributes.

The alc package contains a DHCP object, and has the following members (Table 197).

Table 197 **DHCP Object Members**

Name	Read	Write	Class
htype	X		integer
hlen	X		integer
hops	X		integer
flags	X		integer
ciaddr	X		integer
yiaddr	X		integer
siaddr	X		integer
giaddr	X		integer
chaddr	X		string
sname	X		string
file	X		string
options	X		TLV
sub_ident		X	string
sub_profile_string		X	string
sla_profile_string		X	string

Table 197 **DHCP Object Members (Continued)**

Name	Read	Write	Class
ancp_string		X	string
app_profile_string		X	string
category_map_name		X	string
int_dest_id		X	string

The TLV type provides easy access to the value part of a stream of type-length-value variables, as is the case for the DHCP option field. In the example on 17.3.1, the circuit-ID is accessed as alc.dhcp.options[82][1].

Some DHCP servers do not echo the relay agent option (option 82) when the DHCP message was snooped instead of relayed. For the convenience of the operator, the relay agent option from the request message is returned when alc.dhcp.options[82] is called.

17.3.1 Configuration

As an example, consider script us5.py on Python Scripts, which sets the sub_ident variable based on the circuit ID of three different DSLAMs:

```
import re
import alc
import struct
# ASAM DSLAM circuit ID comes in three flavours:
#       FENT    string          "TLV1: ATM:3/0:100.33"
#       GELT    octet-stream    0x01010000A0A0A0A0000000640022
#       GENT    string          "ASAM11 atm 1/1/01:100.35"
#
# Script sets output ('subscriber') to 'sub-vpi.vci', e.g.: 'sub-
100.33'. circuitid = str(alc.dhcp.options[82][1])
m = re.search(r'(\d+\.\d+)$', circuitid)
if m:
# FENT and GENT
alc.dhcp.sub_ident = "sub-" + m.group()
elif len(circuitid) >= 3:
# GELT
# Note: what byte order does GELT use for the VCI?
# Assume network byte (big endian) order for now. vpi = struct.unpack('B', circuitid
[-3:-2])[0]
vci = struct.unpack('>H', circuitid[-2:])[0]
alc.dhcp.sub_ident = "sub-%d.%d" % (vpi, vci)
```

Configure the URL to this script in a sub-ident-policy as follows:

```
-----------------------------------------------
    sub-ident-policy "DSLAM" create
        description "Parse circuit IDs from different DSLAMs"
        primary

            script-url "ftp://198.51.100.0/py/us5.py"
            no shutdown
        exit
    exit
-----------------------------------------------
```

And attach this sub-ident-policy to the sub-sla-mgmt from a SAP:

```
A:dut-A>config>service>vpls>sap# info
-----------------------------------------------
    dhcp
        description "client side
        lease-populate 50

        no shutdown
    exit
    anti-spoof ip-mac
    sub-sla-mgmt
        sub-ident-policy "DSLAM"
        no shutdown
    exit
-----------------------------------------------
```

Note: DHCP snooping/relaying should be configured properly in order for this to work.

17.3.2 Operator Debugging

Verbose debug output is sent to debug-trace on compile errors, execution errors, execution output and the exported result variables.

```
A:dut-A>config>subscr-mgmt>sub-ident-pol>primary# script-url "ftp://xxx.xxx.xx.xx/
py/
parsefail1.py"
A:dut-A>config>subscr-mgmt>sub-ident-pol>primary# no shutdown

1 2006/07/30 01:17:33.14 UTC MINOR: DEBUG #2001 - Python Compile Error
"Python Compile Error: parsefail1.py
    File "ftp://xxx.xxx.xx.xx/py/parsefail1.py", line 2 def invalid_function():
        ^
IndentationError: expected an indented block
"

A:dut-A>config>subscr-mgmt>sub-ident-pol>primary# script-url "ftp://xxx.xxx.xx.xx/
```

```
py/
dump.py"

2 2006/07/30 01:24:55.50 UTC MINOR: DEBUG #2001 - Python Output
"Python Output: dump.py htype   = 0
hlen    = 0 hops    = 0 flags   = 0
ciaddr  = '0.0.0.0' yiaddr  = '0.0.0.0' siaddr  = '0.0.0.0' giaddr  = '0.0.0.0' chad
dr  = ''
sname   = ''
file    = ''
options = '5\x01\x056\x04\n\x01\x07\n3\x04\x00\x00\x00\xb4\x01\x04\xff\xff\xff
\x00\x1c\x04\n\x02\x02\xffR\x0f\x01\rdut-A|1|1/1/1\xff' "

3 2006/07/30 01:24:55.50 UTC MINOR: DEBUG #2001 - Python Result
"Python Result: dump.py
"

A:dut-A>config>subscr-mgmt>sub-ident-pol>primary# script-url "ftp://xxx.xxx.xx.xx/
py/end- less.py"

4 2006/07/30 01:30:17.27 UTC MINOR: DEBUG #2001 - Python Output
"Python Output: endless.py
"

5 2006/07/30 01:30:17.27 UTC MINOR: DEBUG #2001 - Python Error
"Python Error: endless.py
Traceback (most recent call last):
File "ftp://xxx.xxx.xx.xx/py/
endless.py", line 2, in ? FatalError: script interrupted (timeout)
"
```

Note: All the Python Result events are empty because none of the scripts set any of the output variables.

17.3.3 Python Scripts

Note: The scripts in this section are test scripts and not scripts that the operator would normally use.

dump.py — from alc import dhcp

```
def print_field(key, value):
print '%-8s = %r' % (key, value)

def ipaddr2a(ipaddr):
return '%d.%d.%d.%d' % (
(ipaddr & 0xFF000000) >> 24, (ipaddr & 0x00FF0000) >> 16, (ipaddr & 0x0000FF00) >> 8
, (ipaddr & 0x000000FF))
```

```
print_field('htype',   dhcp.htype) print_field('hlen',    dhcp.hlen) print_field('ho
ps',    dhcp.hops) print_field('flags',   dhcp.flags) print_field('ciaddr',  ipaddr2
a(dhcp.ciaddr)) print_field('yiaddr',  ipaddr2a(dhcp.yiaddr)) print_field('siaddr',
 ipaddr2a(dhcp.siaddr)) print_field('giaddr',  ipaddr2a(dhcp.giaddr)) print_field('c
haddr',  dhcp.chaddr) print_field('sname',   dhcp.sname) print_field('file',    dhcp
.file) print_field('options', str(dhcp.options))
```

us5.py — import re import alc import struct

```
# ASAM DSLAM circuit ID comes in three flavors:
#       FENT   string         "TLV1: ATM:3/0:100.33"
#       GELT   octet-stream   0x01010000A0A0A0A0000000640022
#       GENT   string         "ASAM11 atm 1/1/01:100.35"
#
# Script sets output ('subscriber') to 'sub-vpi.vci', e.g.: 'sub-
100.33'. circuitid = str(alc.dhcp.options[82][1])
m = re.search(r'(\d+\.\d+)$', circuitid)
if m:
# FENT and GENT
alc.dhcp.sub_ident = "sub-" + m.group()
elif len(circuitid) >= 3:
# GELT
# Note: what byte order does GELT use for the VCI?
# Assume network byte (big endian) order for now. vpi = struct.unpack('B', circuitid
[-3:-2])[0]
vci = struct.unpack('>H', circuitid[-2:])[0]
alc.dhcp.sub_ident = "sub-%d.%d" % (vpi, vci)
```

17.3.4 Sample Python Scripts

This section provides examples to show how the script can be used in the context of Enhanced Subscriber Management.

Note: These scripts are included for informational purposes only. The operator must customize the script to match their own network and processes.

17.3.4.1 Example

This script uses the IP address assigned by the DHCP server to derive both *sub_ident* and *sla_profile_string*.

Script:

```
1.  import alc
```

```
2.  yiaddr = alc.dhcp.yiaddr
3.  # Subscriber ID equals full client IP address.
4.  # Note: IP address 10.10.10.10 yields 'sub-168430090'
5.  # and not 'sub-10.10.10.10'
6.  alc.dhcp.sub_ident = 'sub-' + str(yiaddr)
7.  # DHCP server is configured such that the third byte (field) of the IP
8.  # address indicates the session Profile ID.
9.  alc.dhcp.sla_profile_string = 'sp-' + str((yiaddr & 0x0000FF00) >> 8)
```

Explanation:

Line 1 — Imports the library “alc” – Library imports can reside anywhere in the script as long as the items are imported before they are used.

Line 2— Assigns the decimal value of the host’s IP address to a temporary variable “yiaddr”. Line 6: The text “sub_“ followed by yiaddr is assigned to “sub_ident” string.

Line 9— The text “sp-“ followed with the third byte of the IP address is assigned to the “sla- profile” string.

If this script is run, for example, with a DHCP-assigned IP address of:

```
yiaddr = 10.10.0.2
```

The following variables are returned:

```
sub_ident: sub-168427522(hex = A0A00002 = 10.10.0.2)
sla_ident: sp-0
```

17.3.4.2 Example

This script returns the *sub_profile_string* and *sla_profile_string*, which are coded directly in the Option 82 string.

Script:

```
1.  import re
2.  import alc
3.  # option 82 formatted as follows:
4.  # "<subscriber Profile>-<sla-profile>”
5.  ident = str(alc.dhcp.options[82][1])
6.  alc.dhcp.sub_ident = ident
7.  tmp = re.match("(?P<sub>.+)-(?P<sla>.+)", str(ident))
8.  alc.dhcp.sub_profile_string = tmp.group("sub")
9.  alc.dhcp.sla_profile_string = tmp.group("sla")
```

Explanation:

Line 1-2 — Import the libraries “re” and “alc”. Library imports can reside anywhere in the script as long as the items are imported before they are used.

Line 6 — Assigns the full contents of the DHCP Option 82 field to the “sub_ident” variable.

Line 7 — Splits the options 82 string into two parts, separated by “-”.

Line 8 — Assigns the first part of the string to the variable “sub_profile_string”.

Line 9 — Assigns the second part of the string to the variable “sla_profile_string”.

If this script is run, for example, with DHCP option field:

```
options = \x52\x0D\x01\0x0Bmydsl-video
```

The following variables are returned:

```
sub_ident: mydsl-video
sub_profile_string: mydsl
sla_profile_string: video
```

17.3.4.3 Example

This script parses the Option82 “circuit-id” info inserted in the DHCP packet by a DSLAM, and returns the *sub_ident* string.

Script:

```
1.  import re
2.  import alc
3.  import struct
4.  # Nokia 7300 ASAM circuit ID comes in three flavors:
5.  #      FENT   string         "TLV1: ATM:3/0:100.33"
6.  #      GELT   octet-stream   0x01010000A0A0A0A0000000640022
7.  #      GENT   string         "ASAM11 atm 1/1/01:100.35"
8.  #
9.  # Script sets output ('subscriber') to 'sub-vpi.vci',
10. # e.g.: 'sub- 100.33'.
11. circuitid = str(alc.dhcp.options[82][1])
12. m = re.search(r'(\d+\.\d+)$', circuitid)
13. if m:
14.     # FENT and GENT
15.     alc.dhcp.sub_ident = "sub-" + m.group()
16. elif len(circuitid) >= 3:
17.     # GELT
18.     # Note: GELT uses network byte (big endian) order for the VCI
19.     vpi = struct.unpack('B', circuitid[-3:-2])[0]
20.     vci = struct.unpack('>H', circuitid[-2:])[0]
21.     alc.dhcp.sub_ident = "sub-%d.%d" % (vpi, vci)
```

Explanation:

Line 1-2 — Import the libraries “re” and “alc” – Library imports can reside anywhere in the script as long as the items are imported before they are used. Needed if regular expressions are used.

Line 3 — Imports the “struct” library, needed if regular expressions are used.

Line 11 — Assigns the contents of the DHCP Option 82 Circuit-ID field to a temporary variable called “circuitid”.

Line 12 — Parses the “circuitid” and checks for the existence of the regular expression “digit.digit” at the end of the string.

Line 15 — If found, a string containing the text “sub-” followed by these two digits is assigned to the variable “sub-ident”.

Line 16 — If not found, and the length of circuit-id is at least 3.

Line 19 — Parses the “circuitid” and assigns the third-last byte to the temporary variable “vpi”.

Line 20 — Parses the “circuitid” and assigns the last two bytes to the temporary variable “vci”.

Line 21 — Assigns a string containing the text “sub-” followed by vpi and vci to the variable “sub-ident”.

If this script is run, for example, with DHCP option field (assigned by an ASAM with FENT card) containing:

```
options = \x52\x16\x01\x14TLV1: ATM:3/0:100.33
```

(in decimal: 80, 22, 1, 20TLV…)

The following variables are returned:

```
sub_ident: sub-100.33
```

If the above script is run, for example, with a DHCP option field (assigned by an ASAM with GELT card) containing:

```
options = \x52\x10\x01\x0E\x01\x01\x00\x00\xA0\xA0\xA0\xA0\x00\x00\x00\x64 \x00\x22
```

(in decimal: 82, 16, 1, 15, 1, 1, 0, 0, 160, 160, 160, 160, 0, 0, 0, 100, 0, 34; corresponding to VPI 100, VCI 34)

Python returns the following variables:

```
sub_ident: sub-100.34
```

If the above script is run, for example, with a DHCP option field (assigned by an ASAM with GENT card) containing:

```
options = \x52\x1A\x01\x18ASAM11 atm 1/1/01:100.35
```

The following variables are returned:

```
sub_ident: sub-100.35
```

17.3.5 Limitations

'%' operator — While %f is supported, %g and %e are not supported.

Floating Point Arithmetic — The floating point arithmetic precision on the box is less than the precision required by the regression suites of Python. For example, pow(2., 30) equals to

1024.*1024.*1024. until five numbers after the point instead of seven and sqrt(9) equals to 3. for the first seven numbers after the point.

Using the round operator fixes these problems. For example, round(pow(2., 30)) equals round(1024.*1024.*1024.) and round(sqrt(9)) equals 3.

17.4 RADIUS Script Policy Overview

Python scripts for RADIUS AAA packets support manipulation in subscriber management application. This feature is supported on the 7750 SR and 7450 ESS routers in mixed mode.

A Python script can be executed in following cases:

- Before the system sends an access-request packet.
- After the system receives an access-accept packet.
- After the system receives an CoA-request packet.
- Before the system sends an accounting-request packet.

The input of the script is the corresponding original packet; and the output of the packet will be used as the new corresponding packet for further ESM AAA process.

The **radius-script-policy** contains URLs of a primary and an optional secondary Python script, which could be a local CF file path or a FTP URL. The configured **radius-script-policy** could be used in different ESM polices like authentication-policy or radius-accounting-policy.

The following operations are supported within the script:

- Obtain the value of an existing attribute or VSA.
- Modify the value of an existing attribute or VSA.
- Add a new attribute or VSA.
- Remove an existing attribute or VSA.

Note: The following RADIUS attributes or VSA are read-only to Python script:

- Message-Authenticator
- Alc-LI-Action
- Alc-LI-Direction
- Alc-LI-Destination
- Alc-LI-FC
- Alc-LI-Intercept-Id
- Alc-LI-Session-Id

Since Release 12.0R1, users should use a Python policy (instead of a RADIUS script policy) for RADIUS packet manipulation.

17.4.1 Python RADIUS API

The following new Python objects, alc.radius, have the following methods:

- drop(): Drop the packet
- header(): Return the dictionary object that includes RADIUS header information

alc.radius also provides methods to manipulate radius attributes via alc.radius.attributes, which has the following methods:

- get(type): Return the first attribute with specified type as a string
- getTuple(type): The same as above but returns a tuple of strings
- getVSA(vendor, type): Return the first VSA as a string
- getVSATuple(vendor, type): The same as above but returns a tuple of strings
- getExt(type, ext-type): Return a string containing the value of the first Extended Type attribute (RFC 6929, *Remote Authentication Dial In User Service (RADIUS) Protocol Extensions*) matching the specified type and extended type
- getExtTuple(type, ext-type): Return a tuple of strings. For each occurrence of an Extended Type attribute (RFC 6929) matching the specified type and extended type, the attribute value is included as a string entry in the tuple
- getExtVSA(type, vendor, vendor-type): Return a string containing the value of the first Extended Vendor Specific attribute (RFC 6929) matching the specified type, vendor and vendor-type
- getExtVSATuple(type, vendor, vendor-type): Return a tuple of strings. For each occurrence of an Extended Vendor Specific attribute (RFC 6929) matching the specified type, vendor and vendor-type, the attribute value is included as a string entry in the tuple.
- getLongExtTuple(type, ext-type): Return a tuple of strings. For each occurrence of a Long Extended Type attribute (RFC 6929) matching the specified type and extended type, the attribute value is included as a string entry in the tuple. The value of all fragments of the same Long Extended Type attribute are concatenated in a single string entry.
- getLongExtVSATuple(type, vendor, vendor-type): Return a tuple of strings. For each occurrence of a Long Extended Vendor Specific attribute (RFC 6929) matching the specified type, vendor and vendor-type, the attribute value is included as a string entry in the tuple. The value of all fragments of the same Long Extended Vendor Specific attribute are concatenated in a single string entry.
- set(type, value): Set the specified attribute to the value. The value must be either a string or a tuple of strings.

- setVSA(vendor, type, value): Set the specified VSA to the value. The value must be either a string or a tuple of strings.
- setExt(type, ext-type, value): Set the specified Extended Type attribute value (RFC 6929). The value can be specified as a string (single attribute) or a tuple of strings (multiple attributes).
- setExtVSA(type, vendor, vendor-type, value): Set the specified Extended Vendor Specific attribute value (RFC 6929). The value can be specified as a string (single attribute) or a tuple of strings (multiple attributes).
- setLongExt(type, ext-type, value): Set the specified Long Extended Type attribute value (RFC 6929). The value can be specified as a string (single attribute) or a tuple of strings (multiple attributes). When the string value length is too long to fit in a single attribute (more than 251 octets), then the value field is fragmented across multiple attributes. The M-bit (more) is set in all but the last fragment.
- setLongExtVSA(type, vendor, vendor-type, value): Set the specified Long Extended Vendor Specific attribute value (RFC 6929). The value can be specified as a string (single attribute) or a tuple of strings (multiple attributes). When the string value length is too long to fit in a single attribute (more than 246 octets), then the value field is fragmented across multiple attributes. The M-bit (more) is set in all but the last fragment.
- clear(type): Remove the specified attribute
- clearVSA(vendor, type): Remove the specified VSA
- clearExt(type, ext-type): Remove all Extended Type or Long Extended Type attributes (RFC 6929) matching the specified type and extended type
- clearExtVSA(type, vendor, vendor-type): Remove all Extended Vendor Specific or Long Extended Vendor Specific attributes (RFC 6929) matching the specified type, vendor and vendor-type
- isSet(type): Return True if the specified attribute exists, False otherwise
- isVSASet(vendor, type): Return True if the specified VSA exists, False otherwise
- isExtSet(type, ext-type): Return True if an Extended Type or Long Extended Type attribute (RFC 6929) matching the specified type and extended type is present. Else return False
- isExtVSASet(type, vendor, vendor-type): Return True if an Extended Vendor Specific or Long Extended Vendor Specific attribute (RFC 6929) matching the specified type, vendor and vendor-type is present. Else return False.

Extended attribute type, Long Extended attribute type, and Extended Vendor Specific data type are specified in RFC 6929 RADIUS Protocol Extensions.

Both the RADIUS client and RADIUS server must support RFC 6929 RADIUS Protocol Extensions to successfully use the Extended and Long Extended attribute types and the Extended Vendor Specific data type. Therefore it is recommended to check the RADIUS servers for RFC 6929 support and upgrade if needed. As shown in Figure 248, RADIUS Python can be used as a workaround to interwork with a RADIUS server that does not support these extensions. You can use the general purpose [26-6527-242] Alc-Radius-Py VSA of type “octets” for this purpose.

Figure 248 **RADIUS Python: Interworking with RADIUS Server with no RFC 6929 Support**

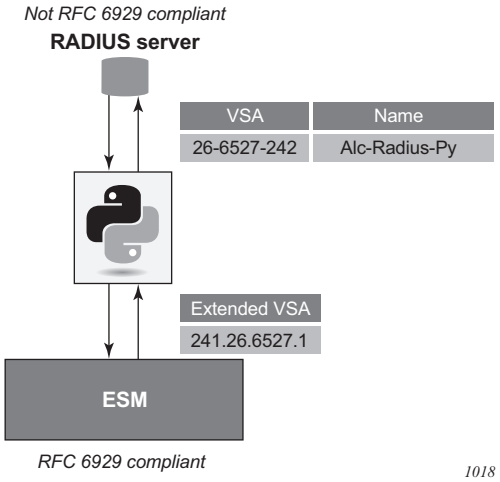

17.4.2 Sample Script

From alc import radius

#1. Get the value of an existing Attribute

```
Username=radius.attributes.get(1)
```

#2. Modify an existing attribute

```
radius.attributes.set(1, 'Tom')
```

#3. Remove an existing attribute

```
radius.attributes.clear(1)
```

#4. Add a new attribute

```
radius.attributes.set(126,”WIFI-operator”)
```

17.5 Python Policy Overview

The Python policy represents a general framework to support all existing and new python features. A Python policy allows users to configure a Python script for a specified ESM packet type (such as DHCP and RADIUS) in a specified direction (ingress or egress). The system executes the configured script when sending or receiving the specified type of packet.

The corresponding original packet is used as input within the script. The user can use the system-provided API to manipulate the input packet (such as add, change, remove option or attribute). The changed packet is used as the output for further ESM processing. With a DHCP transaction cache, the script can return ESM attributes.

Python policies support the following ESM packet types and applications:

- DHCP Transaction Cache on CPM
- DHCPv4 on CPM
- DHCPv6 on CPM
- Diameter on CPM
- GTPv1-C/GTPv2-C
- Python cache on CPM
- RADIUS on CPM
- RADIUS on ISA

There are two types of Python policies:

- Centralized: Runs on the CPM and handles above protocol packets processed on the CPM.
- Distributed: Runs on the ISA and handles packets processed locally on the ISA (such as the DSM). Currently in RADIUS only.

The user must specify the type when creating the Python policy.

The following is an example configuration of a specified group interface. The system will execute cf1:/dhcpv4.py after receiving DHCPv4 discovery and before the system forwards a DHCPv4 request packet.

```
config>python# info
----------------------------------------------
        python-script "dhcpv4" create
            primary-url "cf1:/dhcpv4.py"
            no shutdown
        exit
        python-policy "dhcp" create
```

```
            dhcp discover direction ingress script "dhcpv4"
            dhcp request direction egress script "dhcpv4"
        exit
-----------------------------------------
config>service>vprn>sub-if>grp-if>dhcp# info
-----------------------------------------
                        python-policy "dhcp"
                        server 10.9.9.9
                        lease-populate 100
                        gi-address 192.168.100.1
                        no shutdown
-----------------------------------------
```

17.5.1 Python Policy – RADIUS API

The RADIUS API in Python policy uses the same API of the radius-script-policy.

17.5.2 Python Policy – DHCPv4 API

- The system will provide a Python object for input DHCPv4 packet: alc.dhcpv4.
- alc.dhcpv4 has following attributes to represent the DHCPv4 header fields: Table 198 displays alc.dhcpv4 attributes.

Table 198 **alc.dhcpv4 Attributes**

Class Attrs	DHCPv4 Header Field	Access
alc.dhcpv4.pkt_len	Integer Total length of original DHCPv4 packet (UDP/IP header excluded, including pad option) in bytes	read
alc.dhcpv4.pkt_netlen	Integer Total length of original DHCPv4 packet(UDP/IP header excluded, pad option in the “options” field excluded) in bytes	read
alc.dhcpv4.op	op	read
alc.dhcpv4.htype	htype	read/write
alc.dhcpv4.hlen	hlen	read/write
alc.dhcpv4.hops	hops	read/write
alc.dhcpv4.xid	xid	read
alc.dhcpv4.secs	secs	read/write
alc.dhcpv4.flags	flags	read/write

Table 198 **alc.dhcpv4 Attributes (Continued)**

Class Attrs	DHCPv4 Header Field	Access
alc.dhcpv4.ciaddr	ciaddr	read/write
alc.dhcpv4.yiaddr	yiaddr	read/write
alc.dhcpv4.siaddr	siaddr	read/write
alc.dhcpv4.giaddr	giaddr	read/write
alc.dhcpv4.chaddr	chaddr	read/write
alc.dhcpv4.sname	sname	read/write
alc.dhcpv4.file	file	read/write

All attributes, except alc.dhcpv4.pkt_len and alc.dhcpv4.pkt_netlen, are string with value of the actual bytes in the header.

The following is a list all functions of alc.dhcpv4:

- alc.dhcpv4.drop()
- alc.dhcpv4. getOptionList()
- alc.dhcpv4.pad(min_size=300)
- alc.dhcpv4. get(op_code)
- alc.dhcpv4. set(op_code,valTuple)
- alc.dhcpv4. clear(op_code)
- alc.dhcpv4.get_relayagent() / alc.dhcpv4.set_relayagent(D4OL)
- alc.dhcpv4.get_vendorspecific() / alc.dhcpv4.set_vendorspecific (D4OL)

DHCPv4 allows using sname and file fields to store options. However all DHCPv4 functions will only operate with the “options” field. If a customer wants to manipulate options in the sname/file field, they need to do the parsing work in the script. (extended string.tlvxy method could help here)

- alc.dhcpv4. drop(): The system will drop the result packet
- alc.dhcpv4. getOptionList(): Returns a tuple that includes the option-code of the existing top level DHCPv4 options in the packet.
  - The order of the element in the tuple is as same as the options that appear in the packet.
  - If there are multiple instances of the same option, then each instance is one element in the tuple.
  - Pad option(0) is excluded.

  - End option(255) is included
  - Example: A DHCP discovery packet with msg-type/lease-time/request-addr/parameter-request-list/agent-info/end will return (53,51,50,55,82,255)
- alc.dhcpv4.pad(min_size=300): This function will pad the resulting DHCPv4 packet to the specified min_size with pad option(0) after executing the whole script. Padding will not be added if the result packet is already >=min_size. The default value of min_size is 300. Although not defined in DHCPv4 RFC, many DHCPv4 implementation has a minimal length requirement of 300 bytes. So this function could pad the result packet to the specified min_size.
- alc.dhcpv4. get(op_code): Returns a tuple that includes all instances of the specified top level option as a string. The value of this string is the exact bytes of the option as it appears in packet(excludes option-code and option-len).
  - If the specified option does not exist, then the function will return ()
  - If the certain instance of specified option does not have the value (len=0 or doesn't have len and value part), then the function will return "" for that instance in the tuple.
  - Example: There is an address lease time option(51) in the packet, with value 60, then alc.dhcpv4.get(51) should return: ('\x00\x00\x00\x3c',)
- alc.dhcpv4. set(op-code,valTuple): This function will in fact remove the all existing instances of specified top level option and insert a list of new options. Each element in valTuple is a string, representing one instance of the new option to be inserted; For each new option, the option-code is specified in op-code. The option-len is the length of the element. The rest of option is the element itself.
  - Example: To insert an address lease time option(51) in the packet, with the value 60; use alc.dhcpv4.set(51, ('\x00\x00\x00\x3c',))
- alc.dhcpv4. clear(op-code): This function will remove the all existing instances of specified top level option.
- Although alc.dhcpv4.get() and alc.dhcpv4.set() provide a generic way to manipulate DHCPv4 top level options, but some DHCPv4 options have a complex/hierarchical structure like option82 and option43. To provide a friendly access to these kinds of options, the system provides the following options' specific functions:
  - alc.dhcpv4.get_relayagent() / alc.dhcpv4.set_relayagent(D4OL)
  - alc.dhcpv4.get_vendorspecific() / alc.dhcpv4.set_vendorspecific (D4OL)

  All alc.dhcpv4.get_XXX() will return a data structure:"D4OL" (DHCPv4 Option List)
  - D4OL is a list. Each element in the list represents an instance of that option. For example, if there are 3 option82 in the "options" field of packet, then get_relayagent() will return a list of 3 elements. Each element represents one instance of the option in the packet.
  - Each element in D4OL is a dict (called dict as "D4O" in this example):

- The key of D4O is the sub-option- code. The value is a list of strings of sub-option-value of all instance of the sub-option.

All alc.dhcpv4.set_XX(OPDL) will accept D4OL as the parameter. Remove all existing instances of the corresponding options and then insert the new options represented by specified D4OL.

Examples:

For a packet with an option82 like following

```
Option: (82) Agent Information Option
    Length: 22
    Option (82) Suboption: (1) Agent Circuit ID
        Length: 8
        Agent Circuit ID: 4a616e737656e73
    Option 82 Suboption: (2) Agent Remote ID
        Length: 10
        Agent Remote ID 62617369632364617461
```

The option-data is (hex formatted)
“01:08:4a:61:6e:73:73:65:6e:73:02:0a:62:61:73:69:63:23:64:61:74:61”

The following is an example script:

```
import alc

option82_list=alc.dhcpv4.get_relayagent()
#option82_list will be [{1:['\x4a\x61\x6e:73\x73\x65\x6e\x73',],2:['\x62\x61\x73\x69
\x63\x23\x64\x61\x74\x61',]},]
Option82_list[0][2][0]='basic#video' #change remote-id to 'basic#video'
alc.dhcpv4.set_relayagent(option82_list)#update the option82
```

17.5.3 Python Policy – DHCPv6 API

- The system provides a Python object for input DHCPv6 packet: alc.dhcpv6
- alc.dhcpv6 has following attributes to represent the DHCPv6 header fields:
  Table 199 displays DHCPv6 header fields.

Table 199 **DHCPv6 Header Fields**

Class Attrs	DHCPv6 Header Field	Client/ Server Msg	Relay Msg	Access
alc.dhcpv6.pkt_len	Integer, Total length of original DHCPv4 packet(UDP/IP header excluded) in bytes	✓	✓	Read

Table 199 **DHCPv6 Header Fields (Continued)**

Class Attrs	DHCPv6 Header Field	Client/ Server Msg	Relay Msg	Access
alc.dhcpv6.msg_type	msg-type	✓	✓	Read
alc.dhcpv6.transaction_id	transaction-id	✓		Read
alc.dhcpv6.hop_count	hop-count		✓	read/write
alc.dhcpv6.link_addr	link-address		✓	read/write
alc.dhcpv6.peer_addr	peer-address		✓	read/write

All header fields (as the attribute of alc.dhcpv6 class) are strings(except pkt_len) with exact bytes as it appears in the packet.

If certain attribute does not exist in the given msg-type, for example if the link_attr does not exist in client/server message(C/S msg), then its value should be None.

- The following is a list of all functions in the class:
  - alc.dhcpv6.drop()
  - alc.dhcpv6. getOptionList()
  - alc.dhcpv6. get(op-code)
  - alc.dhcpv6. set(op-code,valTuple)
  - alc.dhcpv6. clear(op-code)
  - alc.dhcpv6.get_iana() / alc.dhcpv6.set_iana(OPDL)
  - alc.dhcpv6.get_iata() / alc.dhcpv6.set_iata(OPDL)
  - alc.dhcpv6.get_vendoropts() / alc.dhcpv6.set_vendoropts(OPDL)
  - alc.dhcpv6.get_iapd() / alc.dhcpv6.set_iapd(OPDL)
  - alc.dhcpv6.get_relaymsg()
  - alc.dhcpv6.set_relaymsg(packet)
- alc.dhcpv6. drop(): The system will drop the resulting packet.
- alc.dhcpv6. getOptionList(): Returns a tuple that includes the option-code of the existing top level DHCPv6 options in the packet. The order of the element in the tuple is as same as the options appear in the packet. If there are multiple instances of same option, then each instance is one element in the tuple. For example:
  - A C/S Msg with Elapsed Time/Client Identifier/IANA/FQDN/Vendor Class/ Option Request will return (8,1,3,39,16,6).
  - A Relay Msg with Relay Message option only will return (9)

- alc.dhcpv6. get(op-code): Returns a tuple that includes all instances of the specified top level option as string The value of this string is the exact bytes of the option as it appears in packet (excludes option-code and option-len). If the specified option does not exist in the input packet, then it will return ().

  Examples:

  - If there is a Status Code option in the packet, status-code 0 and status-msg:”Address Assigned”; then alc.dhcpv6.get(13) should return: (‘\x00\x00Address Assigned’,)

- alc.dhcpv6. set(op-code,valTuple): This function will remove the all existing instances of the specified top level option and insert a list of new options. Each element in valTuple is a string, representing one instance of the new option to be inserted. For each new option, the option-code is specified in op-code, the option-len is the length of the element, reset of option is the element itself.

  - To insert a Status Code options with status-code 0 and status-msg:”Address Assigned”; use alc.dhcpv6.set(13, (‘\x00\x00Address Assigned’,))

- alc.dhcpv6. clear(op-code): This function will remove the all existing instances of specified top level option.

- Although alc.dhcpv6.get() and alc.dhcpv6.set() provide a generic way to manipulate DHCPv6 top level options, but some DHCPv6 options have more complex/hierarchical structure like IA_NA/IA_TA, etc. To provide a friendly access to these kinds of options, the system provides the following options specific functions:

  - alc.dhcpv6.get_iana() / alc.dhcpv6.set_iana(OPDL)
  - alc.dhcpv6.get_iata() / alc.dhcpv6.set_iata(OPDL)
  - alc.dhcpv6.get_vendoropts() / alc.dhcpv6.set_vendoropts(OPDL)
  - alc.dhcpv6.get_iapd() / alc.dhcpv6.set_iapd(OPDL)

  All alc.dhcpv6.get_XXX() will return a data structure:“OPDL” (Option Data structure List)

  - OPDL is a list. Each element in the list represents an instance of that option. For example, if there are 3 IANA in the packet, then get_iana() will return a list of 3 elements, each element represent one instance of IANA option in the packet.
  - Each element in OPDL is a list, referred to as “OPD” in this list.
    - Each element in OPD represent one field in the option(option-code and option-len are not included), the order of the element is as same as the fields appear in the option
    - For field that could be parsed into sub-option by RFC, then the element is a dict, the key of this dict is the sub-option type, if sub-option is one of following supported-sub-option, the value to the key is a sub-option_OPDL represent the list of that specific sub-option

- IAADDR(5)
- Status Code(13)
- IAPREFIX(26)

Else, if the sub-option is **not** one of above, then the value to the key is a list of string of sub-option-data, each string represent one instance of the sub-option.

- The structure of sub-option_OPDL of IAADDR is: [[v6_addr,prefer_lifetime, valid_lifetime,sub-option_OPDL], etc.]
- The structure of sub-option_OPDL of Status Code is: [[status-code,status-msg], etc.]
- The structure of sub-option_OPDL of IAPREFIX is: [[prefer_lifetime,valid_lifetime,prefix-len,v6prefix,sub-option_OPDL],..]
- For the field (by RFC definition) could be parsed into sub-options, but it does not actually exist, then the dict will be empty {}
- For field that cannot be parsed into sub-option by RFC, the element is a string of exact bytes of that field

All alc.dhcpv6.set_XX(OPDL) will accept an OPDL as the parameter. Remove all existing instances of the corresponding options and then insert new options represented by the specified OPDL.

- alc.dhcpv6.get_iana()/alc.dhcpv6.get_iana(OPDL)
- The general OPDL structure for these two functions is: [[IAID_val,T1_val,T2_val, sub-option_dict]]
- The structure of sub-option_dict is:{sub-option-type:sub-option_val}
- If sub-option is supported-sub-option, then sub-option_val is a sub-option_OPDL
- For all other sub-options, the sub-option_val is a list of string of sub-option-data

Examples:

For a packet with an IANA option as in Figure 249:

Figure 249

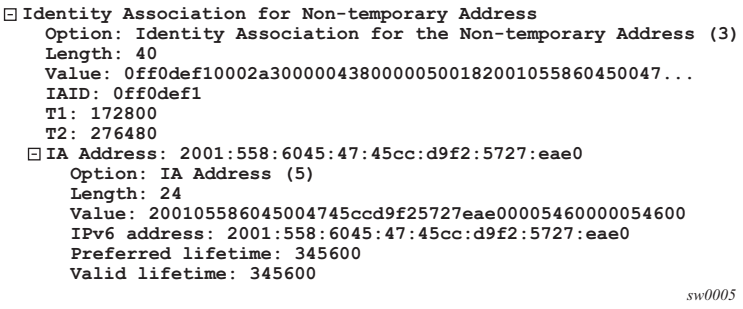

The option-data is (hex formatted)

```
"0f:f0:de:f1:00:02:a3:00:00:04:38:00:00:05:00:18:20:01:05:58:60:45:00:47:45:cc:d9:f2
:57:27:ea:e0:00:05:46:00:00:05:46:00"
```

The following is an example script:

```
import alc

iana_list=alc.dhcpv6.get_iana()
#iana_list will be [['\x0f\xf0\xde\xf1','\x00\x02\xa3\x00','\x00\x04\\x38\x00',{5:[[
'\x20\x01\x05\x58\x60\x45\x00\x47\x45\xcc\xd9\xf2\x57\x27\xea\xe0','\x00\x05\x46\x00
','\x00\x05\x46\x00',{}]]}]]
iana_list[0][1]='\x00\x00\x04\xb0' #change T1 to 1200
alc.dhcpv6.set_iana(iana_list)#update the iana
```

- alc.dhcpv6.get_iata()/alc.dhcpv6.get_iata(OPDL)
  - The general OPDL structure for these two functions is: [[IAID_val, sub-option_dict]]
  - The structure of sub-option_dict is:{sub-option-type:sub-option_val}
  - If sub-option is supported-sub-option, then sub-option_val is a sub-option_OPDL
  - For all other sub-options, the sub-option_val is a list of string of sub-option-data

  Examples: These two function are very similar with IANA, so the examples are skipped here.
- alc.dhcpv6.get_iapd()/alc.dhcpv6.get_iapd(OPDL)
- The general OPDL structure for these two functions is: [[IAID_val,T1_va1,T2_val, sub-option_dict]]
- The structure of sub-option_dict is:{sub-option-type:sub-option_val}
- If sub-option is supported-sub-option, then sub-option_val is a sub-option_OPDL

- For all other sub-options, the sub-option_val is a list of string of sub-option-data

Examples: For a packet with IA_PD as in Figure 250:

Figure 250

```
⊟ Identity Association for Prefix Delegation
    Option: Identity Association for Prefix Delegation (25)
    Length: 41
    Value: 000000010000070800000b40001a001900000e1000015180...
    IAID: 00000001
    T1: 1800
    T2: 2880
  ⊟ IA Prefix
      Option: IA Prefix (26)
      Length: 25
      Value: 00000e10000151803820010db80002000000000000000000...
      Preferred lifetime: 3600
      Preferred lifetime: 86400
      Prefix length: 56
      Prefix address: 2001:db8:2::
```

sw0006

The option-data is (hex formatted)

```
"00:00:00:01:00:00:07:08:00:00:0b:40:00:1a:00:19:00:00:0e:10:00:01:51:80:38:20:01:0d
:b8:00:02:00:00:00:00:00:00:00:00:00:00"
```

Following is an example script:

```
import alc

iapd_list=alc.dhcpv6.get_iapd()
#iapd_list will be [['\x00\x00\x00\x01','\x00\x00\x07\x08','\x00\x00\x0b\x40',{26:[[
'\x00\x00\x0e\x10','\x00\x01\x51\x80','\x38', '\x20\x01\x0d\xb8\x00\x02\x00\x 00\x00
\x00\x00\x00\x00\x00\x00\x00',{}]]}]]
iapd_list[0][1]='\x00\x00\x04\xb0' #change T1 to 1200
alc.dhcpv6.set_iapd(iapd_list)#update the iapd
```

- alc.dhcpv6.get_vendoropts()/alc.dhcpv6.get_vendoropts (OPDL)
  - The general OPDL structure for these two functions is: [[enterpriseid_val, sub-option_dict]]
  - The structure of sub-option_dict is:{sub-option-type:sub-option_val}
  - If sub-option is supported-sub-option, then sub-option_val is a sub-option_OPDL
  - For all other sub-options, the sub-option_val is a list of string of sub-option-data

Examples: For a packet with vendor options as in Figure 251:

Figure 251

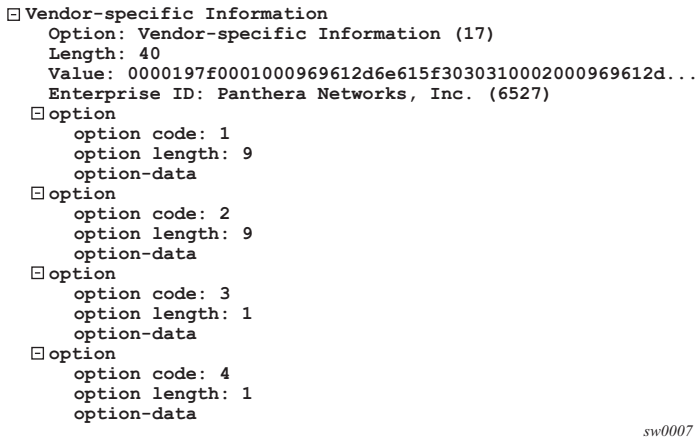

The option-data is (hex formatted)

```
"00:00:19:7f:00:01:00:09:69:61:2d:6e:61:5f:30:30:31:00:02:00:09:69:61:2d:70:64:5f:30
:30:31:00:03:00:01:38:00:04:00:01:40"
```

The following is an example script:

```
import alc

vendoropts_list=alc.dhcpv6.get_vendoropts()
# vendoropts_list will be [['\x00\x00\x19\x7f', {1:['\x69\x61\x2d\x6e\x61\x5f\x30\x3
0\x31'],2:[ '\x69\x61\x2d\x70\x64\x5f\x30\x30\x31'],3:['\x38'], 4:['\x40']}]]
iapd_list[0][1][4][0]='\x60' #change sub-option 4's value to 0x60
alc.dhcpv6.set_vendoropts(vendoropts_list)#update the vendor options
```

- For DHCPv6 relay message, the "Relay Message" option embedded a full DHCPv6 packet and the embedded packet could itself have a "Replay Message" option which embedded another DHCPv6 packet.

  To provide direct access to embedded DHCPv6 packet, the system provides following functions:

  – alc.dhcpv6.get_relaymsg()

  – alc.dhcpv6.set_relaymsg(packet)
- alc.dhcpv6. get_relaymsg(): This function returns a populated alc.dhcpv6 object. which means the returned object was initialized with the DHCPv6 packet embedded in "Relay Message" option as the input.
- alc.dhcpv6. set_relaymsg(packet): This function accepts an alc.dhcpv6 object as a parameter. This object replaces existing "Relay Message" option.

Example-1 script for single relay message:

```
import alc
#input packet is a relay-reply msg
embed_dhcpv6_packet=alc.dhcpv6.get_relaymsg()
iana_list=embed_dhcpv6_packet.get_iana()
iana_list[0][1]='\x00\x00\x04\xb0' #change T1 to 1200
embed_dhcpv6_packet.set_iana(iana_list)#update the iana of the embedded packet
alc.dhcpv6.set_relaymsg(embed_dhcpv6_packet)#update the Relay Message option
```

Example-2 script for double relay message (relay of relay):

```
import alc
#input packet is a relay-reply msg
embed_lv1_packet=alc.dhcpv6.get_relaymsg() #get the level=1 embedded packet
embed_lv2_packet= embed_lv1_packet.get_relaymsg()#get level-
2 packet embedded in level-1 packet
iana_list=embed_dhcpv6_packet.get_iana()
iana_list[0][1]='\x00\x00\x04\xb0' #change T1 to 1200
embed_dhcpv6_packet.set_iana(iana_list)#update the iana
embed_lv1_packet.set_relaymsg(embed_lv2_packet)#update the Relay Message option of l
v1 msg
alc.dhcpv6.set_relaymsg(embed_lv1_packet)#update the Relay Message option of the top
 level msg
```

17.5.4 Python Policy – Diameter API

The alc.diameter Python module provides an API for Diameter message manipulation.

Terminology used in the API description:

- **top-level-AVP** — AVP appearing at the top level in a Diameter message, i.e. not embedded in the Data field of a grouped AVP
- **embedded-AVP** — AVP embedded in the Data field of a grouped AVP. An embedded AVP can be a grouped AVP. This is called nesting.
- **AVP-tuple —** Python tuple with following values:

  (AVP code, vendor id, flags)
  - AVP code : integer, AVP header field
  - vendor id : integer, AVP header field
  - flags : string, AVP header field
- **AVP-value-tuple —** Python tuple with following values:

  (flags, data)
  - flags: string, AVP header field
  - data: string, AVP data field
- **AVP-key-tuple —** Python tuple with following values:

(AVP code, vendor id)
  – AVP code: integer, AVP header field
  – vendor id : integer, AVP header field

- **grouped-AVP-value-tuple —** Python tuple with following values:
  (flags, grouped-AVP-dictionary)
  – flags: string, AVP header field
- **grouped-AVP-dictionary** - Python dictionary with following key:value pairs:
  {AVP-key-tuple : [AVP-value-tuple or grouped-AVP-value-tuple, …], ... }
  – key = AVP-key-tuple
  – value = list of AVP-value-tuples or list of grouped-AVP-value-tuples
- **grouped-AVP-decode-tuple —** Python tuple with following values:
  (AVP-key-tuple, …)
  – tuple of AVP-key-tuples
- **AVP-order-tuple —** Python tuple with following values:
  (AVP-key-tuple, …)
  – tuple of AVP-key-tuples

Table 200 displays attributes available in **alc.diameter** module providing access to the Diameter message header:

Table 200 **Diameter Message Header alc.diameter Attributes**

Attribute	Description	Type	Access
application_id	Diameter message header field	string	Read/Write
code	Diameter message header field	string	Read/Write
end_to_end_id	Diameter message header field	string	Read/Write
flags	Diameter message header field	string	Read/Write
hop_by_hop_id	Diameter message header field	string	Read/Write
msg_length	Diameter message header field. The value is the message length of the original diameter message.	string	Read
version	Diameter message header field	string	Read/Write

Table 201 displays methods available in **alc.diameter** module providing message and AVP manipulation functionality:

Table 201 **Message and AVP Manipulation Functionality alc.diameter Methods**

Method	Description
clear_avps(AVP code, vendor id) AVP code, vendor id = top-level-AVP	Remove all instances of the specified AVP from the message. Applies to top-level-AVP's only. If the specified AVP is not present, no python error is generated. Vendor ID value zero matches top-level-AVP's without Vendor Id field. Return value: None For example: diameter.clear_avps(256, 12645)
drop()	Drop the Diameter message. Packet is consumed at TCP level (ack send). A drop will trigger retransmits on Diameter level. Return value: None For example: diameter.drop()
get_avps(AVP code, vendor id) AVP code, vendor id = top-level-AVP	Returns a list of AVP-value-tuples. Each AVP-value-tuple represents an instance of the specified AVP in the message. Applies to top-level-AVP's only. The position in the list corresponds with the position of the AVP instance in the message at that stage in the script. When executed before any clear or set AVP method, the list order corresponds with the AVP order in the received message. If the specified AVP is a grouped AVP, then the data will contain all the embedded-AVP's. An empty list is returned if the specified AVP is not present. Vendor id value zero matches top-level-AVP's without Vendor Id field. For example: diameter.get_avps(263, 0) [('@', 'bng.nokia.com;1398156449;28')]
get_avps_list()	Returns a list of AVP-tuples. Each AVP-tuple represents an instance of an AVP in the message. Applies to top-level-AVP's only. The position in the list corresponds with the position of the AVP in the message at that stage in the script. When executed before any clear or set AVP method, the list order corresponds with the AVP order in the received message. When multiple instances of an AVP are present in the message, then there will be multiple instances in the list. The Vendor Id has value zero when not present. Grouped AVPs cannot be distinguished from other AVPs in the list. For example: diameter.get_avps_list() [(263, 0, '@'), (264, 0, '@'), (296, 0, '@'), (258, 0, '@'), (416, 0, '@'), (415, 0, '@'), (268, 0, '@'), (55, 0, '@'), (456, 0, '@'), (456, 0, '@'), (456, 0, '@'), (293, 0, '@'), (256, 12645, '\x80')]

Table 201 **Message and AVP Manipulation Functionality alc.diameter Methods (Continued)**

Method	Description
get_grouped_avps(AVP code, vendor id, grouped-AVP-decode-tuple) AVP code, vendor id = top-level-AVP	Returns a list of grouped-AVP-value-tuples with each grouped-AVP-dictionary entry representing an embedded AVP. Each grouped-AVP-value-tuple represents an instance of the specified AVP in the message. Applies to top-level-AVPs of type grouped only. The position in the list corresponds with the position of the grouped AVP instance in the message at that stage in the script. When executed before any clear or set AVP method, the list order corresponds with the AVP order in the received message. If the grouped-AVP-decode-tuple is empty, only the specified top-level-AVP is expanded in a grouped-AVP-value-tuple, with each grouped-AVP-dictionary entry representing an embedded AVP and all dictionary values of type “list of AVP-value-tuples” To expand nested AVPs (grouped AVPs embedded in a grouped AVP), the grouped top-level-AVP and grouped embedded-AVP to expand must be added to the grouped-AVP-decode-tuple. All grouped AVPs in the grouped-AVP-decode-tuple will be expanded in a list of grouped-AVP-value-tuples provided that their embedding AVP is also in the list. The position of the embedded AVPs in the grouped-AVP-dictionary does not correspond with the position in the grouped AVP. If the specified top-level-AVP is not a grouped AVP, then a Python error is generated: “ValueError: malformed diameter message”. For example: To expand the Multiple Services Credit Control (456) grouped top level AVP: diameter.get_grouped_avps(456,0,()) [('@', {(432, 0): [('@', '\x00\x00\x00\x01')], (431, 0): [('@', '\x00\x00\x01\xa4@\x00\x00\x0c\x00\x00\x00d')], (448, 0): [('@', '\x00\x00\x04\xb0')], (268, 0): [('@', '\x00\x00\x07\xd1')]}), ('@', {(432, 0): [('@', '\x00\x00\x00\x02')], (431, 0): [('@', '\x00\x00\x01\xa4@\x00\x00\x0c\x00\x00\x03\x84')], (448, 0): [('@', '\x00\x00\x04\xb0')], (268, 0): [('@', '\x00\x00\x07\xd1')]}), ('@', {(432, 0): [('@', '\x00\x00\x00\x03')], (431, 0): [('@', '\x00\x00\x01\xa4@\x00\x00\x0c\x00\x00\x00<')], (448, 0): [('@', '\x00\x00\x04\xb0')], (268, 0): [('@', '\x00\x00\x07\xd1')]})] To expand the nested Granted-Service-Unit AVP (code 431) in the grouped Multiple Services Credit Control top level AVP (code 456): diameter.get_grouped_avps(456,0,((456,0),(431,0)))

Table 201 **Message and AVP Manipulation Functionality alc.diameter Methods (Continued)**

Method	Description
	[('@', {(432, 0): [('@', '\x00\x00\x00\x01')], (431, 0): [('@', {(420, 0): [('@', '\x00\x00\x00d')]})], (448, 0): [('@', '\x00\x00\x04\xb0')], (268, 0): [('@', '\x00\x00\x07\xd1')]}), ('@', {(432, 0): [('@', '\x00\x00\x00\x02')], (431, 0): [('@', {(420, 0): [('@', '\x00\x00\x03\x84')]})], (448, 0): [('@', '\x00\x00\x04\xb0')], (268, 0): [('@', '\x00\x00\x07\xd1')]}), ('@', {(432, 0): [('@', '\x00\x00\x00\x03')], (431, 0): [('@', {(420, 0): [('@', '\x00\x00\x00<')]})], (448, 0): [('@', '\x00\x00\x04\xb0')], (268, 0): [('@', '\x00\x00\x07\xd1')]})]
set_avps(AVP code, vendor id, list of AVP-value-tuples) AVP code, vendor id = top-level-AVP	Remove all instances of the specified top-level-AVP from the message. For each entry in the AVP-value-tuple list, a top-level-AVP instance is inserted. If the specified vendor id value is zero, then no vendor id field is inserted and setting the Vendor-Specific bit in the flags field of the AVP value tuple will then result in a Python error: “ValueError: no vendor ID but vendor flag set”. If the specified vendor id value is non-zero, then a vendor id field is inserted and not setting the Vendor-Specific bit in the flags field of the AVP value tuple will result in a Python error: “ValueError: vendor ID but vendor flag not set”. Padding between AVPs, AVP length and Diameter message length are adapted accordingly by the system. Return value is None. For example: diameter.set_avps(461,0,[('\x40', 'Pytho.n-1'), ('\x40', 'Python-2')])

Table 201 **Message and AVP Manipulation Functionality alc.diameter Methods (Continued)**

Method	Description
set_fixed_position_avps(AVP-order-tuple)	Put the specified top-level-AVPs at the beginning of the message in the order as specified in the AVP-order-tuple. This method overrides the order of the top-level-AVPs in the resulting Diameter message. If for example the session-id AVP must appear as first in the message, then the corresponding AVP-key-tuple must be included in the first position of the AVP-order-tuple. AVPs not present in the message but specified in the AVP-order-tuple are ignored. AVPs present in the message and not specified in the AVP-order-tuple will be included in the final message after the AVPs listed in the AVP-order-tuple. The order is deterministic but implementation specific. This method can appear at any point in the script. The last call will override the previous one. From a black box viewpoint, this method is executed at the end of the script: the result of the call is not reflected in the list returned by a subsequent get_avp_list() call. Return value: None For example: diameter.set_fixed_position_avps(((263,0), (264,0), (296,0), (268,0)
set_grouped_avps(AVP code, vendor id, list of grouped-AVP-value-tuples) AVP code, vendor id = top-level-AVP	Remove all instances of the specified grouped top-level-AVP from the message. For each entry in the grouped-AVP-value-tuples list, a grouped top-level-AVP instance is inserted. The order of the embedded-AVPs in the grouped AVP cannot be specified. If the specified vendor id value is zero, then no vendor id field is inserted and setting the Vendor-Specific bit in the flags field of the AVP value tuple will then result in a Python error: “ValueError: no vendor ID but vendor flag set”. If the specified vendor id value is non-zero, then a vendor id field is inserted and not setting the Vendor-Specific bit in the flags field of the AVP value tuple will result in a Python error: “ValueError: vendor ID but vendor flag not set”. Padding between AVPs, AVP length and Diameter message length are adapted accordingly. Return value is None. For example: diameter.set_grouped_avps(456,0,[('@', {(432, 0): [('@', '\x00\x00\x00\x01')], (431, 0): [('@', {(420, 0): [('@', '\x00\x00\x00\x2b')]})], (448, 0): [('@', '\x00\x00\x04\xb0')], (268, 0): [('@', '\x00\x00\x07\xd1')]}), ('@', {(432, 0): [('@', '\x00\x00\x00\x03')], (431, 0): [('@', {(420, 0): [('@', '\x00\x00\x00\x53')]})], (448, 0): [('@', '\x00\x00\x04\xb0')], (268, 0): [('@', '\x00\x00\x07\xd1')]})])

To enable Diameter message manipulation using Python, a Python policy must be configured in the diameter-peer-policy. For example:

The Python policy specifies the Python script to use for each Diameter message type received or transmitted on a Diameter peer. In the ingress direction, the Python script is executed when the corresponding packet type is received on the Diameter peer but prior to further processing in the system. In the egress direction, the Python script is executed prior to sending the corresponding packet type on the Diameter peer. For example:

The Python script specifies the location of the script and optional protection mechanism. For example:

```
configure
    python
        python-script "diameter-1" create
            primary-url "ftp://usr:pwd@192.0.2.1/./py/diam-1.py"
            no shutdown
        exit
    exit
```

As an example, the diam-1.py script, clears the M-bit from the Event-Timestamp AVP (code 55):

```
from alc import diameter
avp55=diameter.get_avps(55,0)
diameter.set_avps(55,0,[('\x00',avp55[0][1])])
```

17.5.5 Python Policy – DHCP Transaction Cache API

A DHCP transaction cache (DTC) is a short-lived cache during DHCPv4/v6 transaction. The cache could be used to store user-chosen information or return ESM attributes via a Python script. The DTC's lifetime is only during a single DHCP transaction (for example, only between discovery-offer, request-reply). This includes both alc.dtc.store() data and alc.dtc.setESM() data. DTC is also a transaction specific cache, which means the cached information could only be accessed by the Python script running in same DHCP transaction.

The following are the DTC APIs:

- alc.dtc.derivedId: A string used as a LUDB lookup key
- alc.dtc.store(cache-key, cache-value): Store the value with the specified cache-key in DTC, both key and value are string.
- alc.dtc.retrieve(cache-key): Returns the cached value string according to the specified cache-key, raise exception if specified key does not exist.

- alc.dtc.setESM(ESM-key, value): Sets the specified ESM attribute, which could be used when system creating the ESM host.

Note: Due to the short-live nature of DTC, setESM should be used in final DHCP transaction before system create ESM host., such as DHCPv4 REQUEST-ACK.

- The following is a list of supported ESM-keys and corresponding Python type:
  - alc.dtc.accountingPolicy:str
  - alc.dtc.ancpString:str
  - alc.dtc.appProfileString:str
  - alc.dtc.catMapString:str
  - alc.dtc.defGw:str
  - alc.dtc.dhcpv4GIAddr:str
  - alc.dtc.dhcpv4Pool:str
  - alc.dtc.dhcpv4ServerAddr:str
  - alc.dtc.dhcpv6LinkAddr:str
  - alc.dtc.dhcpv6PreferredLifetime
  - alc.dtc.dhcpv6RebindTimer
  - alc.dtc.dhcpv6RenewTimer
  - alc.dtc.dhcpv6ServerAddr:str
  - alc.dtc.dhcpv6ValidLifetime
  - alc.dtc.intDestId:str
  - alc.dtc.ipAddress:str
  - alc.dtc.ipv4LeaseTime:int
  - alc.dtc.ipv4PrimDns:str
  - alc.dtc.ipv4SecDns:str
  - alc.dtc.ipv6Address:str
  - alc.dtc.ipv6DelegatedPrefix:str
  - alc.dtc.ipv6DelegatedPrefixLength:int
  - alc.dtc.ipv6PrefixPool:str
  - alc.dtc.ipv6PrimDns:str
  - alc.dtc.ipv6SecDns:str
  - alc.dtc.ipv6SlaacPrefix:str
  - alc.dtc.ipv6WanPool:str
  - alc.dtc.msapGroupInterface:str

- alc.dtc.msapPolicy:str
- alc.dtc.msapServiceId:str
- alc.dtc.primNbns:str
- alc.dtc.primNbns:str
- alc.dtc.retailServiceId:str
- alc.dtc.secNbns:str
- alc.dtc.slaProfileString:str
- alc.dtc.subIdent:str
- alc.dtc.subnetMask:str
- alc.dtc.subProfileString:str

17.5.6 Python for PPPoE API

The following are the system-provided Python API PPPoE packets:

- The system provides a Python object for input PPPoE packet: alc.pppoe.
- alc.pppoe class is always available even when Python script is trigger by other type of packet that encapsulated in PPPoE, such as LCP/IPCP/PAP/CHAP, and so on.
- alc.pppoe has the following attributes to represent the PPPoE header fields.

Table 202 displays information about system-provided Python API PPPoE packets.

Table 202 **System Provided Python API PPPoE Packets**

Class Attrs	PPPoE Header Field	Access
alc.pppoe.dest_mac	Str, destination MAC address; the format is “xx:xx:xx:xx:xx:xx”	Read
alc.pppoe.src_mac	Str, source MAC address; the format is “xx:xx:xx:xx:xx:xx”	Read
alc.pppoe.port_id	Str, the port id of the PPPoE session like “1/1/3”	Read
alc.pppoe.vlan_tag	Str, the vlan tag of the PPPoE session like “100” or “100.200”	Read
alc.pppoe.ver	Integer, ver field in PPPoE header	Read
alc.pppoe.type	Integer, type field in PPPoE header	Read

Table 202 **System Provided Python API PPPoE Packets (Continued)**

Class Attrs	PPPoE Header Field	Access
alc.pppoe.code	Integer, code field in PPPoE header	Read
alc.pppoe.session_id	Integer, session_id field in PPPoE header	Read
alc.pppoe.len	Integer, length field in PPPoE header	Read

The following is a list of functions of alic.pppoe:

- alc.pppoe.drop() — The system drops the resulting packet.
- alc.pppoe.getTagList() — The system returns a tuple that includes tag-type of existing PPPoE tags in the packet. The order of the element in the tuple is as same as the tags appear in the packet. If there are multiple instances of same option, then each instance is one element in the tuple.
- alc.pppoe.get(tag-type) — The system returns a tuple of strings, each string represent one instance of the specified tags , the value of this string is exact bytes of the tag_value as it appears in packet(excludes tag_type and tag_len); if tag does not exist, return (); if a specific instance's tag_length=0, return "" for this instance.
- alc.pppoe.set(tag-type,valTuple) — This function removes the all existing instances of specified tags and insert a list of new tags. Each element in valTuple is a string, represent one instance of the new tag to be inserted. For each new option, the type is specified in tag-type, the option-len is the length of the element, reset of option is the element itself.
- alc.pppoe.clear(tag-type) — This function removes the all existing instances of specified tags. In fact this function is a shortcut version of alc.pppoe.set(tag-type,()).
- alc.pppoe.getPPP() — This function returns an alc.ppp object which represent the LCP/IPCP/IP6CP packet encapsulated inside the PPPoE packet. It will return none if there is no such packet.
- alc.pppoe.getPAP() —This function returns an alc.pap object which represents the PAP packet encapsulated inside the PPPoE packet. It will return none if there is no PAP packet;
- alc.pppoe.setPAP(pap_obj) — This function accepts an alc.pap object as the parameter and replace the existing PAP packet that encapsulated in the PPPoE packet. An exception is raised when there is no existing PAP packet.
- alc.pppoe.getCHAP() — This function returns an alc.chap object, which represent the CHAP packet encapsulated inside the PPPoE packet. It will return none if there is no CHAP packet.

- alc.pppoe.setCHAP(chap_obj) — This function accepts an alc.chap object as the parameter and replace the existing CHAP packet that encapsulated in the PPPoE packet. An exception is raised when there is no existing CHAP packet.

17.5.7 Python API for PPP Packet

The system will provide a Python object for PPP packet encapsulated in PPPoE packets: alc.ppp.

- Use alc.pppoe.getPPP() to get the alc.ppp object in the input PPPoE packet. This class only supports following control level protocol:
  - LCP (read-only)
  - IPCP(read-only)
  - IP6CP(read-only)
  - alc.ppp has following attributes to represent the PPP header fields:

Table 203 displays information about PPP header field attributes.

Table 203 **PPP Header Field Attributes**

Class Attrs	PPP Header Field	Access
alc.ppp.protocol	Integer, protocol field in ppp header	Read
alc.ppp.code	Integer, code field in ppp protocol header	Read
alc.ppp.id	Integer, identifier field in ppp protocol header	Read
alc.ppp.len	Integer, length field in ppp protocol header	Read

The following is a list of functions of alc.ppp:

- alc.ppp.getOptionList() — Returns a tuple that includes type of existing PPP options in the packet. The order of the element in the tuple is as same as the options appear in the packet. If there are multiple instances of same option, then each instance is one element in the tuple.
- alc.ppp.get(op-type) — Returns a tuple of strings, each string represent one instance of the specified option , the value of this string is exact bytes of the option_data as it appears in packet (excludes option_type and option_len); if the option does not exist, return (); if a specific instance's option_length=0, return "" for this instance.

17.5.8 Python API for PPP PAP

The system will provide a python object for PAP packet encapsulated in the input PPPoE packets: alc.pap.

Use alc.pppoe.getPAP() to get the alc.pap object in the input PPPoE packet.

To apply changes to PAP packet, alc.pppoe.setPAP(new_alc.pap_obj) must be called.

Table 204 provides PAP header field information. alc.pap has following attributes to represent the PAP header fields:

Table 204 **PAP Header Fields**

Class Attrs	PAP Field	Access
alc.pap.code	Integer, code field in pap protocol header	Read
alc.pap.id	Integer, identifier field in pap protocol header	Read
alc.pap.len	Integer, length field in pap protocol header	Read

The following is a list of functions in the class:

- alc.pap.getCred() — Only apply to Authentication-Request packet. Returns a tuple: (peer-id, password), both are str. If either of them has 0 length, it will be “”; if the system failed to parse the packet, such as the length < 6 bytes or the packet is not an authentication-request, then the system raises an exception. If the packet has a wrong type it returns None (that is, no exception is raised).
- alc.pap.setCred(cred_tuple) — Only apply to Authentication-Request packet. Set the peer-id and password to cred_tuple; cred_tuple is a tuple:(peer-id, password) both are str. “” is allowed as the value for either of them. If the packet is not authentication-request, then system will raise an exception.
- alc.pap.getMsg() — Only apply to Authentication-Ack/Nak packet. Return the message in the packet as a str; return “” if len of message is 0; raise an exception for wrong type of packet (such as an auth-request). If the packet has a wrong type it returns None (that is, no exception is raised).
- alc.pap.setMsg(msg) — Only apply to Authentication-Ack/Nak packet. set the message in the packet to parameter msg; msg is a str; “” is allowed; raise an exception for wrong type of packet (such as an auth-request).

17.5.9 Python API for PPP CHAP

The system will provide a python object for CHAP packet encapsulated in the input PPPoE packets: alc.chap.

Use alc.pppoe.getCHAP() to get the alc.chap object in the input PPPoE packet.

To apply changes to PAP packet, alc.pppoe.setCHAP(new_alc.pap_obj) need to be called.

Table 205 provides CHAP header field information. alc.chap has following attributes to represent the CHAP header fields:

Table 205 **CHAP Header Fields**

Class Attrs	PAP Field	Access
alc.chap.code	Integer, code field in chap protocol header	Read
alc.chap.id	Integer, identifier field in chap protocol header	Read
alc.chap.len	Integer, length field in chap protocol header	Read

The following is a list of functions in the class:

- alc.chap.getCred() — Only apply to challenge/response packet. Returns a tuple: (name,challenge, response), both are str. If either of them has 0 length, it will be “”; if system failed to parse the packet, such as the length < minimal_length or the packet is not challenge/response, then system raise an exception. If the packet has a wrong type it returns None (that is, no exception is raised). The second value in the returned tuple either contains the challenge (in a challenge packet) or the response (in a response packet) value.
- alc.chap.setCred(chap_cred_tuple) — Only apply to challenge/response packet packet. Set the name and challenge to chap_cred_tuple; chap_cred_tuple is a tuple: (name, challenge, response) both are str. “” is allowed as the value for either of them. If packet is not challenge/response, then system will raise an exception. The second value in the supplied tuple either contains the challenge (in a challenge packet) or the response (in a response packet) value.
- alc.chap.getMsg() — Only apply to Success/Failure packet. Return the message in the packet as a str; return “” if len of message is 0; raise an exception for wrong type of packet (such as challenge). If the packet has a wrong type it returns None (that is, no exception is raised).
- alc.chap.setMsg(msg) — Only apply to Success/Failure packet. set the message in the packet to parameter msg; msg is a str; “” is allowed; raise an exception for wrong type of packet(such as challenge).

17.5.10 Python ESM API

The system provides a python object: alc.esm, to provide flexible LUDB lookup and return ESM attributes directly from script.

alc.esm support both DHCPv4/v6 and PPPoE hosts.

alc.esm has following attribute and function:

- alc.esm.derivedId: A string used as LUDB lookup key
- alc.esm.set(ESM-key, value): Set the value of specified ESM attribute, it will be used for ESM host creation.

alc.esm.derivedId should be set for all types of ingress packets that system used to access LUDB; for example PPPoE PADI and PADR, DHCP discovery and request.

The following is a list of supported ESM-key and its corresponding Python type:

- alc.esm.accountingPolicy:str
- alc.esm.ancpString:str
- alc.esm.appProfileString:str
- alc.esm.catMapString:str
- alc.esm.defGw:str
- alc.esm.dhcpv4GIAddr:str
- alc.esm.dhcpv4Pool:str
- alc.esm.dhcpv4ServerAddr:str
- alc.esm.dhcpv6LinkAddr:str
- alc.esm.dhcpv6ServerAddr:str
- alc.esm.intDestId:str
- alc.esm.ipAddress:str
- alc.esm.ipv4LeaseTime:int
- alc.esm.ipv4PrimDns:str
- alc.esm.ipv4SecDns:str
- alc.esm.ipv6Address:str
- alc.esm.dhcpv6PreferredLifetime
- alc.esm. dhcpv6RebindTimer
- alc.esm. dhcpv6RenewTimer
- alc.esm. dhcpv6ValidLifetime
- alc.esm.ipv6DelegatedPrefix:str

- alc.esm.ipv6DelegatedPrefixLength:int
- alc.esm.ipv6PrefixPool:str
- alc.esm.ipv6PrimDns:str
- alc.esm.ipv6SecDns:str
- alc.esm.ipv6SlaacPrefix:str
- alc.esm.ipv6WanPool:str
- alc.esm.msapGroupInterface:str
- alc.esm.msapPolicy:str
- alc.esm.msapServiceId:str
- alc.esm.primNbns:str
- alc.esm.retailServiceId:str
- alc.esm.secNbns:str
- alc.esm.slaProfileString:str
- alc.esm.spiSharingGroupId:int
- alc.esm.subIdent:str
- alc.esm.subProfileString:str
- alc.esm.subnetMask:str
- alc.dtc.dhcpv6PreferredLifetime
- alc.dtc.dhcpv6RebindTimer
- alc.dtc.dhcpv6RenewTimer
- alc.dtc.dhcpv6ValidLifetime

The following are PPPoE-specific keys:

- alc.esm.padoDelay: int
- alc.esm.pppAuthProtocol: str, {‘pap’|’chap’|’pref-chap’|’pref-pap’}
- alc.esm.pppMTU: int

With PPPoE, alc.esm.set() could only be called ingress discovery (PADI, PADR) and authentication (PAP Authenticate, CHAP Response) packets. If it is called multiple times, the first one wins except the derived-id, while the alc.esm.padoDelay only applies when the system receives PADI. alc.esm.pppAuthProtocol and alc.esm.pppMTU can only be set when a PADR is received.

alc.esm.set() will be blocked during LCP/IPCP/IP6CP phase.

The system will raise an exception if alc.esm.set() is called at wrong time. Setting the ESM keys that do not apply to the host type are ignored; for example, setting PPP(oE) keys for DHCP hosts is ignored.

The value returned by alc.esm.set() take highest precedence than other sources like CLI/LUDB.

For alc.esm.set(ESM-key, value)

- PPPoE can only be called for ingress discovery (PADI, PADR) and authentication (PAP Authenticate, CHAP Response) packets. An exception is raised otherwise (for example, wrong time, including egress discovery PADO and PADS, and authentication PAP Ack/Nak & CHAP Success/Failure).
- The PPPoE key alc.esm.padoDelay can only be set when PADI is received (otherwise, it is too late).
- PPP keys alc.esm.pppAuthProtocol and alc.esm.pppMTU can only be set when PADR is received. It has no effect when PADI is received.
- Setting ESM-keys for a PPPoE host which do not apply is ignored (that is, no exception is raised).
- Setting PPP(oE) keys for a DHCP host is ignored.

17.5.11 Python Cache Support

Python cache support allows information to be shared across different run times of the same python script or even different python scripts in a programmatic way. It essentially provides a central memory cache and a set of APIs which let the user store and retrieve strings. For example, a DHCP python-script could store a DHCP option into cache and later a RADIUS python-script could retrieve stored string and add it into access-request.

Each cached entry in the cache is a tuple of (**key**, **val**). **key** is used as entry id. **val** is the string to be cached. Both **key** and **val** are strings. The max length of the key is 512 bytes. Future more, the combine length of key+val is limited by the configured value of **entry-size** *size* command in the python-policy.

The Python cache can be enabled per python-policy. Each python policy has its own cache memory which script in other python-policy cannot access. This also implies that the key of a cached entry in different a python policy could overlap.

The user can also specify the max number cache entry per python policy the command **max-entries** command. System has a global limit for python cache memory of 256MB.

The cached entries could be made persistent by saving it to CF card. This can be enabled with the **persistence** command in the python policy.

Note: From memory consumption point of view, with MCS enabled, each cached entry will have a corresponding MCS record, so each entry will consume twice amount of memory.

The system also supports syncing the python cache across chassis with MCS. This can be configured per python policy with the **mcs-peer** command in the python policy.

Each cached entry has a remaining lifetime. If it decreases over time, the system will remove the cached entry if its remaining lifetime is 0. The remaining lifetime can be changed using a system- provided API. The initial lifetime of a newly created cache entry is 600 seconds.

The following are the python cache APIs in a module alc.cache:

- alc.cache.save(val,key): Saves the val identified by the key into the cache. If there is an existing cache entry with same key, then it will be overwritten with the new val. An exception will be raised if the save failed (for example, due to exceeding the max number of entries).
- alc.cache.retrieve(key): Returns the stored entry's val identified by the key. A KeyError exception will be raised if the specified entry does not exist.
- alc.cache.clear(key): Removes the cached entry identified by the key. Raise KeyError exception if the specified entry does not exist.
- cache.get_lifetime(key): The system will return a integer as seconds of remaining lifetime of the specified entry. It will return none if the specified entry does not exist. An exception will be raised for any other error.
- cache.set_lifetime(key,new_lifetime): The new_lifetime value is an integer. The system will set the remaining lifetime of the specified entry to the number of seconds of the new_lifetime. An exception will be raised for any error including specified entry does not exist. If the new_lifetime>=max_lifetime(configurable using the **max-entry-life** command in the python policy), then the system will set the actual lifetime to the max_lifetime.

17.5.12 Applying a Python Policy

The following is a list of places that a Python policy could be applied:

- Under capture SAP — Apply to the DHCPv4/v6 packets sent/received on the capture SAP
- Under group-interface — Apply to DHCPv4/v6 packets sent/received on the group-interface

- Under subscriber-interface — Apply to DHCPv4 packets on the retail subscriber interface
- In the radius-server-policy — Apply to the RADIUS packets sent/received to/from the RADIUS servers configured in the radius-server-policy
- In the radius-proxy-server — Apply to the RADIUS packets on the client side of proxy
- In the diameter-peer-policy — Apply to the Diameter packets send/received on the Diameter peers configured in the policy

17.5.13 Python Script Protection

In order to protect the Python script from unintended changes, the SR-OS supports a new Python script file format:SRPY. SRPY includes a key based hash(HMAC) of the original script content. When the system loads a script with SRPY format, a hash will be computed by using a configured key and script content. The result hash will be compared to the embedded hash. If it is the same, then this script is considered valid. Otherwise, the system will abort with a warning message.

Users can configure **protection hmac-sha256 key** *key* within a Python script. To mandate, all configured scripts must be in SRPY format.

The system provides a tools command (**tools perform python-script protect**) to convert a Python script into SRPY format.

17.6 Tips and Tricks

- Use xrange() instead of range().
- Avoid too many string operations. The following scripts provide the same output:

```
# This script takes 2.5 seconds. s = ""
for c in 'A'*50000:
s += str(ord(c)) + ', ' print '[' + s[:-2] + ']'

# This script takes 0.1 seconds. print map(ord, 'A'*50000)
```

17.7 Python Command Reference

17.7.1 Command Hierarchies

- Python Policy Commands
- Python Script Commands
- Services Commands
- Show Commands
- Tools Commands
- Debug Commands
- Clear Commands

17.7.1.1 Python Policy Commands

config
- — **python**
  - — **python-policy** *name* [**create**] [**wlan-gw-group** *wlan-gw-group-id*] [**nat-group** *nat-group-id*]
  - — **no python-policy** *name*
    - — **cache** [**create**]
    - — **no cache**
      - — **entry-size** *size*
      - — **no entry-size**
      - — **max-entries** *count*
      - — **no max-entries**
      - — **max-entry-lifetime** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
      - — **no max-entry-lifetime**
      - — **mcs-peer** *ip-address* **sync-tag** [*sync-tag*]
      - — **no mcs-peer**
      - — **minimum-lifetimes**
        - — **high-availability** *seconds*
        - — **no high-availability**
        - — **multi-chassis-redundancy** *seconds*
        - — **no multi-chassis-redundancy**
        - — **persistence** *seconds*
        - — **no persistence**
      - — [**no**] **persistence**
      - — [**no**] **shutdown**
    - — **description** *description-string*
    - — **no description**
    - — **dhcp** *type* **direction** {**ingress** | **egress**} **script** *script*
    - — **no dhcp** *type* **direction** {**ingress** | **egress**}

```
            — dhcp6 type direction {ingress | egress} script script
            — no dhcp6 type direction {ingress | egress}
            — diameter diameter type direction {ingress |egress} script [script]
            — no diameter type direction {ingress | egress}
            — gtpv1-c type direction {ingress| egress} script [script]
            — no gtpv1-c type direction {ingress| egress}
            — gtpv2-c type direction {ingress| egress} script [script]
            — no gtpv2-c type direction {ingress| egress}
            — pppoe type direction {ingress | egress} script name
            — no pppoe type direction {ingress | egress}
            — radius type direction {ingress | egress} script script
            — no radius type direction {ingress | egress}
            — [no] syslog script name
            — vsd script [32 chars max]
            — no vsd

config
    — aaa
        — isa-radius-policy
            — python-policy policy-name
            — no python-policy
```

17.7.1.2 Python Script Commands

```
config
    — python
        — python-script name [create]
        — no python-script name
            — action-on-fail {drop | passthrough}
            — no action-on-fail
            — description description-string
            — no description
            — primary-url url
            — no primary-url
            — protection none
            — protection hmac-sha256 key key [hash|hash2]
            — no protection
            — secondary-url url
            — no secondary-url
            — [no] shutdown
            — [no] tertiary-url url
            — no tertiary-url

config
    — aaa
        — radius-server-policy policy-name [create]
        — no radius-server-policy policy-name
            — python-policy name
            — no python-policy

config
```

```
— router router-name
    — radius-proxy
        — server server-name [create] [purpose {[accounting] [authentication]} [wlan-
            gw-group wlan-gw-group-id]
        — no server server-name
            — python-policy name
            — no python-policy

config
    — system
        — persistence
            — python
                — description description-string
                — no description
                — location cflash-id
                — no location

config
    — redundancy
        — multi-chassis
            — peer
                — sync
                    — [no] python
```

17.7.1.3 Services Commands

Refer to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 2 Services and EVPN Guide: VLL, VPLS, PBB, and EVPN* and to the *7450 ESS, 7750 SR, 7950 XRS, and VSR Layer 3 Services Guide: IES and VPRN* for information on command syntax and CLI descriptions.

```
config
    — service
        — vpls service-id
        — no vpls service-id
            — sap sap-id [create] [capture-sap]
            — no sap sap-id
                — dhcp-python-policy policy-name
                — no dhcp-python-policy
                — dhcp6-python-policy policy-name
                — no dhcp6-python-policy
                — pppoe-python-policy policy-name
                — no pppoe-python-policy

config
    — service
        — ies service-id [create]
        — no ies service-id
        — vprn service-id [create]
        — no vprn service-id
```

```
            — subscriber-interface ip-int-name
            — no subscriber-interface ip-int-name
                — dhcp
                    — python-policy name
                    — no python-policy
                — group-interface ip-int-name
                — no group-interface ip-int-name
                    — dhcp
                        — python-policy name
                        — no python-policy
                    — ipv6
                        — dhcp6
                            — python-policy name
                            — no python-policy
                    — pppoe
                        — [no] python-policy policy-name

config
    — service
        — vprn service-id [create]
        — no vprn service-id
            — radius-proxy
                — server server-name [create] [purpose {[accounting] [authentication]}
                    [wlan-gw-group wlan-gw-group-id]
                — no server server-name
                    — python-policy name
                    — no python-policy
```

17.7.1.4 Show Commands

```
show
    — python
        — python-policy
        — python-policy policy-name [association]|
        — python-script
        — python-script script-name [association | source-in-use]
        — python-script script-name statistics nat-group nat-group-id
        — python-script script-name statistics wlan-gw-group wlan-gw-group-id
```

17.7.1.5 Tools Commands

```
tools
    — dump
        — python
            — python-policy name cache
            — python-policy name cache hex-key hex-string
            — python-policy name cache string-key [string-key]
    — perform
```

— **python-policy** *name* **cache hex-key** *hex-string* **set-lifetime** [*lifetime*]
— **python-policy** *name* **cache** *string-key* [512 chars max] **set-lifetime** [*lifetime*]
— **python-script protect input** *file-url* **hmac-sha256 key** *secret-key* **output** *file-url*
— **python-script reload** *name*

17.7.1.6 Debug Commands

debug
— **python**
— [**no**] **python-script** *name*
— **script-all-info**
— [**no**] **script-compile-error**
— [**no**] **script-export-variables**
— [**no**] **script-output**
— [**no**] **script-output-on-error**
— [**no**] **script-runtime-error**

17.7.1.7 Clear Commands

clear
— **python**
— **python-policy** *name* **cache**
— **python-policy** *name* **cache hex-key** *hex-string*
— **python-policy** *name* **cache string-key** [512 chars max]

17.7.2 Command Descriptions

17.7.2.1 Python Configuration Commands

17.7.2.1.1 Global Commands

description

Syntax **description** *description-string*
no description

Context config>python>py-policy
config>python>python-script

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes the description string from the context.

Parameters *description-string* — A text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, etc.), the entire string must be enclosed within double quotes.

shutdown

Syntax [**no**] **shutdown**

Context config>python>py-policy>cache
config>python>python-script

Description Shutting down a Python script triggers the system to load and compile the script from the configured location(s). Since the system supports three locations, the primary, secondary and tertiary, the system will try to load the Python script in that order.

Shutting down a Python script will disable the Python script and cause the corresponding packet to pass through without any modification.

The **no** form of this command enables the cache or policy script.

Default no shutdown

python

Syntax **python**

Context config

Description This command enables the context to configure Python parameters.

17.7.2.1.2 Python Policy Commands

python-policy

Syntax **python-policy** *name* [**create**] [**wlan-gw-group** *wlan-gw-group-id*] [**nat-group** *nat-group-id*]

no python-policy *name*

Context config>python

Description This command creates a new Python policy or enables an existing Python policy configuration context.

There are two types of Python policies: centralized and distributed. A centralized Python policy runs on a CPM, while a distributed Python policy runs on an ISA. With the distributed Python policy, a **wlan-gw-group** *wlan-gw-group-id* or a **nat-group** *nat-group-id* command must specified.

The **no** form of this command removes the Python policy from the configuration.

Parameters *name* — Specifies the Python policy name up to 32 characters.

create — This keyword is required when first creating the Python policy. Once the context is created, it is possible to navigate into the context without the **create** keyword

wlan-gw-group *wlan-gw-group-id* — Specifies the ID of the WLAN GW group that the distributed python-policy installs

nat-group *nat-group-id* — Specifies the ID of the NAT group that the distributed python-policy installs

python-policy

Syntax **python-policy** *policy-name*
no python-policy

Context config>aaa>isa-radius-policy

Description This command specifies the Python policy for the ISA RADIUS proxy server. This is the python policy for RADIUS packets to/from the client.

The **no** form of this command removes the Python policy from the configuration.

Parameters *name* — Specifies the Python policy name up to 32 characters

create — This keyword is required when first creating the Python policy. Once the context is created, it is possible to navigate into the context without the **create** keyword.

cache

Syntax **cache** [**create**]
no cache

Context config>python>py-policy

Description This command enables the context to configure the limits of the caching API inside the Python scripts.

The **no** form of this command removes the configured cache parameters from the configuration.

Parameters **create** — This keyword is required when first creating the Python policy. Once the context is created, it is possible to navigate into the context without the **create** keyword.

entry-size

Syntax **entry-size** *size*
no entry-size

Context config>python>py-policy>cache

Description This command configures the maximum size of the data structure that can be stored in a single Python cache entry which includes both a value and key.

When requesting to store a data structure, the size of the serialized object is compared with the value specified. If larger, the object will not be stored and Python will return exception.

The **no** form of this command reverts to the default.

Default entry-size 256

Parameters *size* — Configures the maximum accepted size of a single cache entry

Values 32 to 2048

max-entries

Syntax **max-entries** *count*
no max-entries

Context config>python>py-pol>cache

Description This command configures the maximum number of Python cache entries that can be stored in the cache of this Python policy.

If the limit has been reached, a Python exception will be thrown when requested to store another data structure.

The **no** form of this command reverts to the default.

Default max-entries 128000

Parameters *count* — Specifies the maximum number of cache entries allowed

Values 1 to 1000000

max-entry-lifetime

Syntax **max-entry-lifetime** [**days** *days*] [**hrs** *hours*] [**min** *minutes*] [**sec** *seconds*]
no max-entry-lifetime

Context config>python>py-pol>cache

Description This command configures the maximum allowed lifetime for each entry of the Python cache of this Python policy.

When adding data to the Python cache the lifetime of the given object must always be specified. If the specified lifetime is bigger than the configured value, then the value of the **max-entry-lifetime** will be used instead of the lifetime that was specified.

The **no** form of this command reverts to the default.

Default max-entry-lifetime days 1

Parameters **days** *days* — Specifies the maximum lifetime that can be set on a cache entry in days

Values 0 to 7

Default 1

hrs *hours* — Specifies the maximum lifetime that can be set on a cache entry in hours

Values 0 to 23

min *minutes* — Specifies the maximum lifetime that can be set on a cache entry in minutes

Values 0 to 59

sec *seconds* — Specifies the maximum lifetime that can be set on a cache entry in seconds

Values 0 to 59

mcs-peer

Syntax **mcs-peer** *ip-address* **sync-tag** [*sync-tag*]
no mcs-peer

Context config>python>py-pol>cache

Description This command specifies the MCS peer's address and sync-tag for syncing the cached entries of the python-policy. The sync-tag must be match on both chassis.

The **no** form of this command reverts to the default.

Parameters *ip-address* — Specifies the IPv4 address of the MCS peer

sync-tag — Specifies the tag for sync up to 32 characters max

minimum-lifetimes

Syntax **minimum-lifetimes**

Context config>python>py-pol>cache

Description This command enables the context to configure minimum-lifetime of Python cache information.

high-availability

Syntax **high-availability** *seconds*
no high-availability

Context config>python>py-pol>cache>minimum-lifetimes

Description This command specifies the minimum lifetime of an entry that it could be synced across CPM.

The **no** form of this command reverts to the default.

Parameters *seconds* — Specifies the minimal lifetime in seconds.

Values 1 to 600

multi-chassis-redundancy

Syntax **multi-chassis-redundancy** *seconds*
no multi-chassis-redundancy

Context config>python>py-pol>cache>minimum-lifetimes

Description This command specifies the minimum lifetime for a cache entry to be synchronized with the MCS peer.

The **no** form of this command reverts to the default.

Parameters *seconds* — Specifies the multi-chassis redundancy time in seconds

Values 1 to 600

persistence

Syntax **persistence** *seconds*
no persistence

Context config>python>py-pol>cache>minimum-lifetimes

Description This command configures the minimum lifetime for a cache entry to be made persistent.

The **no** form of this command reverts to the default.

Parameters *persistence* — The minimum lifetime in seconds

Values 1 to 600

persistence

Syntax [**no**] **persistence**

Context config>python>py-pol>cache

Description This command enables persistency support for the cached entries of the python-policy.

The **no** form of this command reverts to the default.

dhcp

Syntax **dhcp** *type* **direction** {**ingress** | **egress**} **script** *script*
no dhcp *type* **direction** {**ingress** | **egress**}

Context config>python>py-policy

Description This command specifies the Python script for the specified DHCPv4 packet type in the specified direction.

Multiple **dhcp** command configurations are allowed in the same Python policy.

The **no** form of this command reverts to the default.

Parameters *type* — Specifies the message type of the event

Values discover, offer, request, decline, ack, nak, release, inform, force-renew, lease-query, lease-unassigned, lease-unknown, lease-active

direction {**ingress** | **egress**} — Specifies whether the packet is being received by the system or being sent by the system

script *script* — Specifies the name of the Python script up to 32 characters, that will be used to handle the specified message

dhcp6

Syntax **dhcp6** *type* **direction** {**ingress** | **egress**} **script** *script*
no dhcp6 *type* **direction** {**ingress** | **egress**}

Context config>python>py-policy

Description This command specifies the Python script for the specified DHCPv6 packet type in the specified direction.

Multiple **dhcps** command configurations are allowed in the same Python policy.

The **no** form of this command reverts to the default.

Parameters *type* — Specifies the message type of the event

Values solicit, advertise, request, confirm, renew, rebind, reply, release, decline, reconfigure, info-request, relay-forward, relay-reply

direction {**ingress** | **egress**} — Specifies whether the event is incoming or outgoing

script *script* — Specifies the name of the Python script up to 32 characters, that will be used to handle the specified message

diameter

Syntax **diameter** *type* **direction** {**ingress** | **egress**} **script** *script*
no diameter *type* **direction** {**ingress** | **egress**}

Context config>python>py-policy

Description This command specifies the Python script to use for the specified Diameter message type in the specified direction.

Multiple diameter command configurations are allowed in the same Python policy.

The **no** form of this command reverts to the default.

Parameters *type* — Specifies the message type

Message type	Application	Direction
aaa – AA Answer	Nasreq	ingress
aar – AA Request	Nasreq	egress
asa – Abort Session Answer	Gx, Gy	egress
asr – Abort Session Request	Gx, Gy	ingress
cca – Credit Control Answer	Gx, Gy	ingress
ccr – Credit Control Request	Gx, Gy	egress
cea – Capabilities Exchange Answer	Base	ingress
cer – Capabilities Exchange Request	Base	egress

dpa – Disconnect Peer Answer	Base	ingress/egress
dpr – Disconnect Peer Request	Base	ingress/egress
dwa – Device Watchdog Answer	Base	ingress/egress
dwr – Device Watchdog Request	Base	ingress/egress
raa – Re-Authentication Answer	Gx, Gy	egress
rar – Re-Authentication Request	Gx, Gy	ingress

direction {**ingress** | **egress**} — Specifies if the message is incoming or outgoing

script *script* — Specifies the name of the Python script up to 32 characters, that will be used to handle the specified message

gtpv1-c

Syntax **gtpv1-c** *type* **direction** {**ingress** | **egress**} **script** [*script*]
no gtpv1-c *type* **direction** {**ingress** | **egress**}

Context config>python>py-policy

Description This command configures a Python script for the specified GTPv1-C message type in the specified direction.

The **no** form of this command reverts to the default.

Parameters *type* — Specifies the message type

Values echo-request, echo-response, version-not-supported, create-pdp-context-request, create-pdp-context-response, delete-pdp-context-request, delete-pdp-context-response, error-indication

direction {**ingress** | **egress**} — Specifies if the message is incoming or outgoing

script *script* — Specifies the name of the Python script up to 32 characters, that will be used to handle the specified message

gtpv2-c

Syntax **gtpv2-c** *type* **direction** {**ingress** | **egress**} **script** [*script*]
no gtpv2-c *type* **direction** {**ingress** | **egress**}

Context config>python>py-policy

Description This command configures a Python script for the specified GTPv2-C message type in the specified direction.

The **no** form of this command reverts to the default.

Parameters *type* — Specifies the message type

Values echo-request, echo-response, version-not-supported, create-session-request, create-session-response, delete-session-request, delete-session-response, delete-bearer-request, delete-bearer-response, modify-bearer-request, modify-bearer-response, release-access-bearers-request, release-access-bearers-response, downlink-data-notification, downlink-data-notification-ack

direction {**ingress** | **egress**} — Specifies if the message is incoming or outgoing

script *script* — Specifies the name of the Python script up to 32 characters, that will be used to handle the specified message

pppoe

Syntax **pppoe type direction** {**ingress** | **egress**} **script** *name*
no pppoe type direction {**ingress** | **egress**}

Context config>python>py-policy

Description This command specifies the python-script for the specified PPPoE message type in the specified direction. Multiple **pppoe** command configuration are allowed in the same Python policy.

The **no** form of this command reverts to the default.

Parameters **type** — Specifies the message type

Values session-lcp, session-pap, session-chap, session-ipcp, session-ip6cp, pado, padi, padr, pads, padt

direction {**ingress** | **egress**} — Specifies whether the event is incoming or outgoing. The system will only invoke the configured script for the specified packet type in the specified direction

script *script* — Specifies the name of the Python script up to 32 characters, that will be used to handle the specified message

radius

Syntax **radius** *type* **direction** {**ingress** | **egress**} **script** *script*
no radius *type* **direction** {**ingress** | **egress**}

Context config>python>py-policy

Description This command specifies the Python script for the specified RADIUS packet type in the specified direction.

Multiple **radius** command configurations are allowed in the same Python policy.

The **no** form of this command reverts to the default.

Parameters *type* — Specifies the message type of the event

Values access-request, access-accept, access-reject, accounting-request, accounting-response, access-challenge, disconnect-request, change-of-authorization-request

direction {**ingress** | **egress**} — Specifies whether the event is incoming or outgoing

script — Specifies the name of the Python script up to 32 characters, that will be used to handle the specified message

syslog

Syntax [**no**] **script script**

Context config>python>py-policy

Description This command enables Python script to process syslog related messages and events.

The **no** form of this command disables the Python script to process syslog related messages and events

Parameters *script* — Specifies the name of the Python script up to 32 characters, which will be used to handle the specified message

vsd

Syntax **vsd script** *script*
no vsd

Context config>python>py-policy

Description This command configures scripts to handle VSD messages.

The **no** form of this command removes the script from the Python policy.

Parameters *script* — Specifies the name of the Python script up to 32 characters, that will be used to handle the specified message

17.7.2.1.3 Python Script Commands

python-script

Syntax **python-script** *name* [**create**]

no python-script *name*

Context config>python

Description This command enables the context to configure Python scripts to modify messages of different protocols.

The **no** form of this command removes the Python script name from the configuration.

Parameters **name** — Specifies the name of this Python script policy

create — This keyword is required when first creating the Python script. Once the context is created, it is possible to navigate into the context without the **create** keyword.

action-on-fail

Syntax **action-on-fail** {**drop** | **passthrough**}
no action-on-fail

Context config>python>py-script

Description This command specifies the action taken when Python fails to modify the given message.

The **no** form of this command reverts to the default.

Default action-on-fail drop

Parameters **drop** — Specifies that the packet will be dropped

passthrough — Specifies that he packet will be sent out without any modifications

primary-url

Syntax **primary-url** *url*
no primary-url

Context config>python>py-script

Description This command specifies the location of the primary Python script. The system supports three locations for each Python script. Users can store the script file on either a local CF card or an FTP server.

The **no** form of this command removes the URL.

Parameters **url** — Specifies the primary URL of the Python script up to 180 characters, either a local CF card url or a FTP server URL

protection

Syntax **protection none**
protection hmac-sha256 key *key* [**hash**|**hash2**]
no protection

Context config>python>py-script

Description This command specifies the format of the Python script file(s) in this python-script. Unintentional changing of Python script file could be prevented by using protected format.

The **no** form of this command equals to **protection none**.

Parameters **none** — Indicates the Python script is stored in plain-text, without any mechanism in place to ensure the integrity nor the confidentiality of the content of the Python script

hmac-sha256 — Indicates the first line of the Python script must consist of the hash value obtained by hashing the rest of the Python script using the hmac-sha256 hashing algorithm given the key specified in tmnxPythonScriptProtectionKey

key *key* — The specified key along with original Python script file content are used to compute the hash. The computed hash will be compared to the hash in the Python script file. If there is no match, then system will fail to load the script.

hash — Specifies the key is entered in an encrypted form. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified

hash2 — Specifies the key is entered in a more complex encrypted form that involves more variables than the key value alone, meaning that the **hash2** encrypted variable cannot be copied and pasted. If the **hash** or **hash2** parameter is not used, the key is assumed to be in an unencrypted, clear text form. For security, all keys are stored in encrypted form in the configuration file with the **hash** or **hash2** parameter specified.

secondary-url

Syntax **secondary-url** *url*
no secondary-url

Context config>python>py-script

Description This command specifies the location of secondary Python script. The system supports three locations for each Python-script. Users can store scripts file on either a local CF card or a FTP server.

The **no** form of this command removes the URL.

Parameters **url** — Specifies the secondary URL of the Python script up to 180 characters, either a local CF card URL or a FTP server URL

tertiary-url

Syntax **tertiary-url** *url*
no tertiary-url

Context config>python>py-script

Description This command specifies the location of tertiary Python script. The system supports three locations for each Python-script. Users can store scripts file on either a local CF card or a FTP server.

The **no** form of this command removes the URL.

Parameters **url** — Specifies the tertiary URL of the Python script up to 180 characters, either a local CF card URL or a FTP server URL

17.7.2.1.4 Services Commands

dhcp-python-policy

Syntax **dhcp-python-policy** *policy-name*
no dhcp-python-policy

Context config>service>vpls>sap

Description This command specified the Python policy for DHCPv4 packets sent/received on the capture SAP.

The **no** form of this command removes the policy name from the configuration.

Parameters *policy name* — Specifies an existing Python policy name

dhcp6-python-policy

Syntax **dhcp6-python-policy** *policy-name*
no dhcp6-python-policy

Context config>service>vpls>sap

Description This command specified the Python policy for DHCPv6 packets sent/received on the capture SAP.

The **no** form of this command removes the policy name from the configuration.

Parameters *policy name* — Specifies an existing Python policy name

pppoe-python-policy

Syntax **pppoe-python-policy** *policy-name*
no pppoe-python-policy

Context config>service>vpls>sap

Description This command specified the Python policy for PPPoE packets sent/received on the capture SAP.

The **no** form of this command removes the policy name from the configuration.

Parameters *policy-name* — Specifies an existing Python policy name

python-policy

Syntax **python-policy** *name*
no python-policy

Context config>service>vprn>sub-if>grp-if>ipv4>dhcp4
config>service>ies>sub-if>grp-if>ipv4>dhcp4

Description This command specified the python-policy for DHCPv4 packets sent/received on the group interface.

The **no** form of this command removes the policy name from the configuration.

Parameters *name* — Specifies an existing Python policy name

python-policy

Syntax **python-policy** *name*
no python-policy

Context config>service>vprn>sub-if>grp-if>ipv6>dhcp6
config>service>ies>sub-if>grp-if>ipv6>dhcp6

Description This command specified the python-policy for DHCPv6 packets sent/received on the group interface.

The **no** form of this command removes the policy name from the configuration.

Parameters *name* — Specifies an existing Python policy name

python-policy

Syntax **python-policy** *policy-name*
no python-policy

Context config>service>ies>sub-if>grp-if>pppoe
config>service>vprn>sub-if>grp-if>pppoe

Description This command specifies the Python policy for PPPoE packets sent/received on the group interface.

The **no** form of this command removes the policy name from the configuration.

Parameters *policy-name* — Specifies an existing Python policy name

python-policy

Syntax **python-policy** *name*
no python-policy

Context config>service>vprn>sub-if>dhcp
config>service>ies>sub-if>dhcp

Description This command specified the Python policy for DHCPv4 packets sent/received on the retail subscriber interface.

The **no** form of this command removes the policy name from the configuration.

Parameters *name* — Specifies the name of the Python policy

python-policy

Syntax **python-policy** *name*
no python-policy

Context config>aaa>radius-srv-plcy

Description This command specifies the python-policy for RADIUS packets to/from the RADIUS servers defined in the specified **radius-server-policy**.

The **no** form of this command removes the policy name from the configuration.

Parameters *name* — Specifies the name of the Python policy

python-policy

Syntax **python-policy** *name*
no python-policy

Context config>router>radius-proxy>server
config>service>vprn>radius-proxy>server

Description This command specifies the python-policy for RADIUS packets sent/received on the client side of the RADIUS proxy server.

This command supports RADIUS proxy on both CPMs and ISAs.

The **no** form of this command removes the policy name from the configuration.

Parameters *name* — Specifies the name of the Python policy

persistence

Syntax [**no**] **persistence**

Context config>system

Description This command enables the context to configure persistence parameters on the system.

The persistence feature enables state on information learned through DHCP snooping across reboots to be retained. This information includes data such as the IP address and MAC binding information, lease-length information, and ingress SAP information (required for VPLS snooping to identify the ingress interface).

If persistence is enabled when there are no DHCP **relay** or **snooping** commands enabled, an empty file is created.

The **no** form of this command reverts to the default.

python-policy-cache

Syntax **python-policy-cache**

Context config>system>persistence

Description This command enables the context to configure Python policy cache persistency parameters.

python

Syntax [**no**] **python**

Context config>redundancy>multi-chassis>peer>sync

Description This command enables syncing of python-policy cached entries to the peer.

Use the **mcs-peer** command in the Python policy to enable syncing for a specific Python policy.

The **no** form of this command reverts to the default.

location

Syntax **location** *cflash-id*
no location

Context config>system>persistence>python

Description This command instructs the system where to write the file. The name of the file is: dhcp-persistence.db. On boot the system scans the file systems looking for dhcp-persistence.db, if it finds it starts to load it.

In the subscriber management context, the location specifies the flash device on a CPMCFM card where the data for handling subscriber management persistency is stored.

The **no** form of this command returns the system to the default. If there is a change in file location while persistence is running, a new file will be written on the new flash, and then the old file will be removed.

The **no** form of this command reverts to the default.

17.7.2.1.5 Tools Commands

python-policy

Syntax **python-policy** *name* **cache**
python-policy *name* **cache hex-key** *hex-string*
python-policy *name* **cache string-key** [*string-key*]

Context tools>dump>python

Description This command dumps all cached entries or a specified entry of a specified Python policy.

The DDP key in the output is the python cache persistency record key.

Note: The DDP Key in the output could be used for the **tools>dump>persistence>python** command.

Parameters *name* — Specifies the name of the Python policy

string-key — Specifies the key of the entry to be updated in ASCII strong format

hex-key — Specifies the key of the entry to be updated in hex string format

python-policy

Syntax **python-policy** *name* **cache hex-key** *hex-string* **set-lifetime** [*lifetime*]
python-policy *name* **cache** *string-key* [512 chars max] **set-lifetime** [*lifetime*]

Context tools>perform

Description This command set the lifetime of a specified python cache entry.

Parameters *name* — Specifies the name of the Python policy

string-key — Specifies the key of the entry to be updated in ASCII strong format

hex-key *hex-string* — Specifies the key of the entry to be updated in hex string format

set-lifetime — Specifies the new lifetime of the entry

Values 1 to 2147483647

python-script protect

Syntax **python-script protect input** *file-url* **hmac-sha256 key** *secret-key* **output** *file-url*

Context tools>perform

Description This command converts a normal (unprotected) Python script file into an SRPY format with specified key.

Parameters **input** *file-url* — Specifies the URL of the input script file

key *secret-key* — Specifies the key used to compute the hash

output *file-url* — Specifies the URL of the output script file

python-script reload

Syntax **python-script reload** *name*

Context tools>perform

Description This command will try to reload/recompile the primary/secondary/tertiary scripts in the specified Python script in order. The system will use the first script that comes up.

Parameters *name* — Specifies the name of the python-script to be reloaded

17.7.2.2 Show Commands

Note: The command outputs in the following section are examples only; actual displays may differ depending on supported functionality and user configuration.

python

Syntax **python**

Context show

Description This command enables the context to display Python information.

python-policy

Syntax **python-policy**
python-policy *policy-name* [**association**]|

Context show>python

Description This command displays information about the currently configured Python policy.

The system will display a list of currently configured Python policy names if no parameter is specified.

Parameters *policy-name* — Specifies the Python policy name to display

association — Displays the associations of the specified Python policy

Output The following output is an example of DHCP Python policy information

Sample Output

```
show python python-policy "dhcp"
===============================================================================
Python policy "dhcp"
===============================================================================
Description   : (Not Specified)
-------------------------------------------------------------------------------
Messages
```

```
-------------------------------------------------------------------------------
Type                                Dir     Script
-------------------------------------------------------------------------------
dhcpDiscover                        egress  dhcpv4
dhcpRequest                         egress  dhcpv4
dhcpAck                             ingress dtc
-------------------------------------------------------------------------------
No. of Messages: 3
-------------------------------------------------------------------------------
===============================================================================


show python python-policy "dhcp" association
===============================================================================
Python Policy Association
===============================================================================
Location
-------------------------------------------------------------------------------
Service: 500, GrpIf g1, dhcp
-------------------------------------------------------------------------------
No. of Python policy association: 1
===============================================================================
```

python-script

Syntax **python-script**
python-script *script-name* [**association** | **source-in-use**

Context show>python

Description This command displays information about the currently configured Python script.

The system will display a list of currently configured Python script names if no parameter is specified.

Parameters *script-name* — Specifies the Python script name to display information

association — Displays the associations of the specified Python script

source-in-use — Displays the Python source code in use

Output The following output is an example of Python script information

Sample Output

```
show python python-script "dhcpv4"
===============================================================================
Python script "dhcpv4"
===============================================================================
Description   : (Not Specified)
Admin state   : inService
Oper state    : inService
Action on fail: drop
Protection    : none
```

```
Primary URL   : cf1:/dhcpv4.py
Secondary URL : (Not Specified)
Tertiary URL  : (Not Specified)
Active URL    : primary
Last changed  : 01/26/2014 05:02:10
===============================================================================


show python python-script "dhcpv4" association
===============================================================================
Python Script Association
===============================================================================
Policy                          Type                                  Dir
-------------------------------------------------------------------------------
dhcp                            dhcpDiscover                          egress
dhcp                            dhcpRequest                           egress
-------------------------------------------------------------------------------
No. of Python script association: 2
===============================================================================
```

17.7.2.3 Debug Commands

python-script

Syntax **python-script** *script-name*

Context debug>python

Description This command enters the debug context the specified Python script

Parameters *policy-name* — Specifies the Python script name

script-all-info

Syntax **script-all-info**

Context debug>python>py-script

Description This command enables the script-compile-error, script-export-variables, script-output, script-output-on-error, and script-runtime-error functionalities.

script-compile-error

Syntax [**no**] **script-compile-error**

Context debug>python>py-script

Description This command sends the traceback of the compile error to the logger. The traceback contains detailed information about where and why the compilation fails. The compilation takes place when the CLI user changes the admin state of the Python script from shutdown to no-shutdown.

script-export-variables

Syntax [**no**] **script-export-variables**

Context debug>python>py-script

Description This command sends the output variables of the Python script to the logger when the script ran successfully.

script-output

Syntax [**no**] **script-output**

Context debug>python>py-script

Description This command sends the output (such as from 'print' statements) of the Python script to the logger.

script-output-on-error

Syntax [**no**] **script-output-on-error**

Context debug>python>py-script

Description This command sends the output (such as traceback data) of the Python script to the logger, but only when the script fails.

script-runtime-error

Syntax [**no**] **script-runtime-error**

Context debug>python>py-script

Description This command generates log information when detecting a script runtime error.

17.7.2.4 Clear Commands

python-policy

Syntax **python-policy** *name* **cache**
python-policy *name* **cache hex-key** *hex-string*
python-policy *name* **cache** *string-key* [512 chars max]

Context clear>python

Description This command clears Python policy data.

Parameters *name* — Specifies the name of the Python policy

string-key — Specifies the key of the entry to be updated in ASCII strong format

hex-key — Specifies the key of the entry to be updated in hex string format

17.8 Python RADIUS Command Reference

17.8.1 Command Hierarchies

```
configure
    — aaa
        — radius-script-policy policy-name [create]
        — no radius-script-policy policy-name
            — action-on-fail {drop | passthrough}
            — no action-on-fail
            — description description
            — no description
            — primary
                — script-url primary-script-url
                — no script-url
                — [no] shutdown
            — secondary
                — script-url secondary-script-url
                — no script-url
                — [no] shutdown

configure
    — aaa
        — l2tp-accounting-policy policy-name [create]
        — no l2tp-accounting-policy policy-name
            — request-script-policy policy-name
            — no request-script-policy
        — radius-server-policy policy-name [create]
        — no radius-server-policy policy-name
            — accept-script-policy policy-name
            — no accept-script-policy
            — request-script-policy policy-name
            — no request-script-policy

configure
    — subscriber-mgmt
        — authentication-policy name [create]
        — no authentication-policy name
            — accept-script-policy policy-name
            — no accept-script-policy
            — coa-script-policy policy-name
            — no coa-script-policy
            — request-script-policy policy-name
            — no request-script-policy
        — radius-accounting-policy name [create]
        — no radius-accounting-policy name
            — acct-request-script-policy policy-name
            — no acct-request-script-policy
```

17.8.2 Command Descriptions

17.8.2.1 Generic Commands

description

Syntax **description** *description-string*
no description

Context config>aaa>radius-script-policy

Description This command creates a text description stored in the configuration file for a configuration context.

The **description** command associates a text string with a configuration context to help identify the context in the configuration file.

The **no** form of this command removes any description string from the context.

Parameters *description-string* — Specifies a text string describing the entity. Allowed values are any string up to 80 characters long composed of printable, 7-bit ASCII characters excluding double quotes. If the string contains special characters (#, $, spaces, and so on), the entire string must be enclosed within double quotes.

shutdown

Syntax [**no**] **shutdown**

Context config>aaa>radius-scr-plcy>primary
config>aaa>radius-scr-plcy>secondary

Description The **shutdown** command administratively disables the entity. When disabled, an entity does not change, reset, or remove any configuration settings or statistics. Many entities must be explicitly enabled using the **no shutdown** command.

The **shutdown** command administratively disables an entity. The operational state of the entity is disabled as well as the operational state of any entities contained within. Many objects must be shut down before they can be deleted.

Unlike other commands and parameters where the default state is not indicated in the configuration file, **shutdown** and **no shutdown** are always indicated in system generated configuration files.

The **no** form of this command puts an entity into the administratively enabled state.

Default no shutdown

17.8.2.2 Script Commands

radius-script-policy

Syntax **radius-script-policy** *policy-name* [**create**]
no radius-script-policy *policy-name*

Context config>aaa

Description This command configures a RADIUS script policy.

The **no** form of this command removes the scrip policy from the configuration.

Parameters *policy-name* — Configures Python scripts to modify RADIUS messages.

create — This keyword is required when first creating the configuration context. Once the context is created, it is possible to navigate into the context without the **create** keyword.

action-on-fail

Syntax **action-on-fail** {**drop** | **passthrough**}
no action-on-fail

Context config>aaa>radius-scr-plcy

Description specifies the action taken when Python fails to modify the RADIUS message.

The **no** form of this command reverts to the default.

Default action-on-fail drop

Parameters **drop** — Specifies that the packet will be dropped.

passthrough — Specifies that the packet will be sent out without any modifications.

primary

Syntax **primary**

Context config>aaa>radius-scr-plcy

Description This command enables the context to configure a primary script.

script-url

Syntax **script-url** *primary-script-url*
no script-url

Context config>aaa>radius-scr-plcy>primary

Description This command configures the URL of the primary script.

The **no** form of this command removes the URL from the configuration.

Parameters *primary-script-url* — Specifies the URL of the secondary script to change RADIUS attributes of the RADIUS message.

secondary

Syntax **secondary**

Context config>aaa>radius-scr-plcy

Description This command enables the context to configure a secondary script.

script-url

Syntax **script-url** *secondary-script-url*
no script-url

Context config>aaa>radius-scr-plcy>secondary

Description Specifies the URL of the secondary script to change RADIUS attributes of the RADIUS message.

The **no** form of this command removes the URL from the configuration.

Parameters *secondary-script-url* — Specifies the URL of the secondary script to change RADIUS attributes of the RADIUS message.

l2tp-accounting-policy

Syntax **l2tp-accounting-policy** *policy-name* [**create**]
no l2tp-accounting-policy *policy-name*

Context config>aaa

Description This command configures an L2TP accounting policy.

The **no** form of this command removes the policy-name from the configuration.

Parameters *policy-name* — Specifies a policy name.

create — This keyword is required when first creating the configuration context. Once the context is created, it is possible to navigate into the context without the **create** keyword.

request-script-policy

Syntax **request-script-policy** *script-policy*
no request-script-policy

Context config>aaa>l2tp-accounting-policy
config>aaa>radius-srv-plcy

Description This command configures a Python script policy to modify Access-Request.

The **no** form of this command removes the policy-name from the configuration.

Parameters *script-policy* — Specifies a the RADIUS script policy used to change the RADIUS attributes of the outgoing Access-Request messages.

radius-server-policy

Syntax **radius-server-policy** *policy-name* [**create**]
no radius-server-policy *policy-name*

Context config>aaa

Description This command configures the RADIUS script policy used to change the RADIUS attributes of the outgoing Access-Accept messages.

The **no** form of this command removes the policy-name from the configuration.

Parameters *policy-name* — Specifies a the RADIUS script policy used to change the RADIUS attributes of the outgoing Access-Request messages.

create — This keyword is required when first creating the configuration context. Once the context is created, it is possible to navigate into the context without the **create** keyword.

authentication-policy

Syntax **authentication-policy** *name* [**create**]
no authentication-policy

Context config>subscr-mgmt

Description This command creates the context to configure RADIUS server parameters for session authentication.

The **no** form of this command removes the RADIUS server configuration for session authentication.

RADIUS servers can be configured for three different applications:

1. For authentication of dynamic Triple Play subscriber sessions, under **config>subscr-mgmt>authentication-plcy**
2. For 802.1x port authentication, under **config>system>security>dot1x>radius-plcy**
3. For CLI login users, under **config>system>radius**

Parameters *name* — Specifies the name of the profile. The string is case sensitive and limited to 32 ASCII 7-bit printable characters with no spaces.

accept-script-policy

Syntax **accept-script-policy** *policy-name*
no accept-script-policy

Context config>aaa>radius-srv-plcy

Description This command configures the RADIUS script policy used to change the RADIUS attributes of the outgoing Access-Accept messages.

The **no** form of this command removes the policy name from the configuration.

Parameters *policy-name* — Specifies the Python script policy to modify Access-Accept.

radius-accounting-policy

Syntax **radius-accounting-policy** *name* [**create**]
no radius-accounting-policy

Context config>subscr-mgmt

Description This command specifies a subscriber RADIUS based accounting policy.

The **no** form of this command removes the policy name from the configuration.

Parameters *name* — Specifies the name of the policy up to 32 ASCII 7-bit printable characters.

coa-script-policy

Syntax **coa-script-policy** *policy-name*

no coa-script-policy

Context config>subscr-mgmt>auth-plcy

Description This command configures the RADIUS script policy used to change the RADIUS attributes of the Change-of-Authorization messages.

The **no** form of this command removes the policy name from the configuration.

Parameters *policy-name* — Specifies the Python script policy to modify the Change-of-Authorization messages.

request-script-policy

Syntax **request-script-policy** *policy-name*
no request-script-policy

Context config>subscr-mgmt>auth-plcy

Description This command configures the RADIUS script policy used to change the RADIUS attributes of the outgoing Access-Request messages.

The **no** form of this command removes the policy name from the configuration.

Parameters *policy-name* — Specifies the Python script policy to modify Access-Request messages.

acct-request-script-policy

Syntax **acct-request-script-policy** *policy-name*
no acct-request-script-policy

Context config>subscr-mgmt>acct-plcy

Description This command configures the Python script policy to modify Accounting-Request messages.

The **no** form of this command removes the policy name from the configuration.

Parameters *policy-name* — Specifies the Python script policy to modify Accounting-Request messages.

18 Standards and Protocol Support

Note: The information presented is subject to change without notice.

Nokia assumes no responsibility for inaccuracies contained herein.

Access Node Control Protocol (ANCP)

draft-ietf-ancp-protocol-02, *Protocol for Access Node Control Mechanism in Broadband Networks*

RFC 5851, *Framework and Requirements for an Access Node Control Mechanism in Broadband Multi-Service Networks*

Application Assurance (AA)

3GPP Release 12 (ADC rules over Gx interfaces)

RFC 3507, *Internet Content Adaptation Protocol (ICAP)*

Asynchronous Transfer Mode (ATM)

AF-ILMI-0065.000, *Integrated Local Management Interface (ILMI) Version 4.0*

AF-PHY-0086.001, *Inverse Multiplexing for ATM (IMA) Specification Version 1.1*

AF-TM-0121.000, *Traffic Management Specification Version 4.1*

AF-TM-0150.00, *Addendum to Traffic Management v4.1 optional minimum desired cell rate indication for UBR*

GR-1113-CORE, *Asynchronous Transfer Mode (ATM) and ATM Adaptation Layer (AAL) Protocols Generic Requirements, Issue 1*

GR-1248-CORE, *Generic Requirements for Operations of ATM Network Elements (NEs), Issue 3*

ITU-T I.432.1, *B-ISDN user-network interface - Physical layer specification: General characteristics (02/99)*

ITU-T I.610, *B-ISDN operation and maintenance principles and functions (11/95)*

RFC 1626, *Default IP MTU for use over ATM AAL5*

RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5*

Bidirectional Forwarding Detection (BFD)

RFC 5880, *Bidirectional Forwarding Detection (BFD)*

RFC 5881, *Bidirectional Forwarding Detection (BFD) IPv4 and IPv6 (Single Hop)*

RFC 5883, *Bidirectional Forwarding Detection (BFD) for Multihop Paths*

RFC 7130, *Bidirectional Forwarding Detection (BFD) on Link Aggregation Group (LAG) Interfaces*

Border Gateway Protocol (BGP)

draft-hares-idr-update-attrib-low-bits-fix-01, *Update Attribute Flag Low Bits Clarification*

draft-ietf-idr-add-paths-guidelines-08, *Best Practices for Advertisement of Multiple Paths in IBGP*

draft-ietf-idr-best-external-03, *Advertisement of the best external route in BGP*

draft-ietf-idr-bgp-flowspec-oid-03, *Revised Validation Procedure for BGP Flow Specifications*

draft-ietf-idr-bgp-gr-notification-01, *Notification Message support for BGP Graceful Restart*

draft-ietf-idr-bgp-optimal-route-reflection-10, *BGP Optimal Route Reflection (BGP-ORR)*

draft-ietf-idr-error-handling-03, *Revised Error Handling for BGP UPDATE Messages*

draft-ietf-idr-flowspec-interfaceset-03, *Applying BGP flowspec rules on a specific interface set*

draft-ietf-idr-flowspec-path-redirect-05, *Flowspec Indirection-id Redirect* (localised ID)

draft-ietf-idr-flowspec-redirect-ip-02, *BGP Flow-Spec Redirect to IP Action*

draft-ietf-idr-link-bandwidth-03, *BGP Link Bandwidth Extended Community*

draft-ietf-sidr-origin-validation-signaling-04, *BGP Prefix Origin Validation State Extended Community*

draft-uttaro-idr-bgp-persistence-03, *Support for Long-lived BGP Graceful Restart*

RFC 1772, *Application of the Border Gateway Protocol in the Internet*

RFC 1997, *BGP Communities Attribute*

RFC 2385, *Protection of BGP Sessions via the TCP MD5 Signature Option*

RFC 2439, *BGP Route Flap Damping*

RFC 2545, *Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing*

RFC 2858, *Multiprotocol Extensions for BGP-4*

RFC 2918, *Route Refresh Capability for BGP-4*

RFC 3107, *Carrying Label Information in BGP-4*

RFC 3392, *Capabilities Advertisement with BGP-4*

RFC 4271, *A Border Gateway Protocol 4 (BGP-4)*

RFC 4360, *BGP Extended Communities Attribute*

RFC 4364, *BGP/MPLS IP Virtual Private Networks (VPNs)*

RFC 4456, *BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)*

RFC 4486, *Subcodes for BGP Cease Notification Message*

RFC 4659, *BGP-MPLS IP Virtual Private Network (VPN) Extension for IPv6 VPN*

RFC 4684, *Constrained Route Distribution for Border Gateway Protocol/ MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)*

RFC 4724, *Graceful Restart Mechanism for BGP* (helper mode)

RFC 4760, *Multiprotocol Extensions for BGP-4*

RFC 4798, *Connecting IPv6 Islands over IPv4 MPLS Using IPv6 Provider Edge Routers (6PE)*

RFC 4893, *BGP Support for Four-octet AS Number Space*

RFC 5004, *Avoid BGP Best Path Transitions from One External to Another*

RFC 5065, *Autonomous System Confederations for BGP*

RFC 5291, *Outbound Route Filtering Capability for BGP-4*

RFC 5396, *Textual Representation of Autonomous System (AS) Numbers* (asplain)

RFC 5549, *Advertising IPv4 Network Layer Reachability Information with an IPv6 Next Hop*

RFC 5575, *Dissemination of Flow Specification Rules*

RFC 5668, *4-Octet AS Specific BGP Extended Community*

RFC 6286, *Autonomous-System-Wide Unique BGP Identifier for BGP-4*

RFC 6810, *The Resource Public Key Infrastructure (RPKI) to Router Protocol*

RFC 6811, *Prefix Origin Validation*

RFC 6996, *Autonomous System (AS) Reservation for Private Use*

RFC 7311, *The Accumulated IGP Metric Attribute for BGP*

RFC 7607, *Codification of AS 0 Processing*

RFC 7674, *Clarification of the Flowspec Redirect Extended Community*

RFC 7752, *North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP*

RFC 7854, *BGP Monitoring Protocol (BMP)*

RFC 7911, *Advertisement of Multiple Paths in BGP*

RFC 7999, *BLACKHOLE Community*

RFC 8092, *BGP Large Communities Attribute*

Circuit Emulation

RFC 4553, *Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)*

RFC 5086, *Structure-Aware Time Division Multiplexed (TDM) Circuit Emulation Service over Packet Switched Network (CESoPSN)*

RFC 5287, *Control Protocol Extensions for the Setup of Time-Division Multiplexing (TDM) Pseudowires in MPLS Networks*

Ethernet

IEEE 802.1AB, *Station and Media Access Control Connectivity Discovery*
IEEE 802.1ad, *Provider Bridges*
IEEE 802.1ag, *Connectivity Fault Management*
IEEE 802.1ah, *Provider Backbone Bridges*
IEEE 802.1ak, *Multiple Registration Protocol*
IEEE 802.1aq, *Shortest Path Bridging*
IEEE 802.1ax, *Link Aggregation*
IEEE 802.1D, *MAC Bridges*
IEEE 802.1p, *Traffic Class Expediting*
IEEE 802.1Q, *Virtual LANs*
IEEE 802.1s, *Multiple Spanning Trees*
IEEE 802.1w, *Rapid Reconfiguration of Spanning Tree*
IEEE 802.1X, *Port Based Network Access Control*
IEEE 802.3ab, *1000BASE-T*
IEEE 802.3ac, *VLAN Tag*
IEEE 802.3ad, *Link Aggregation*
IEEE 802.3ae, *10 Gb/s Ethernet*
IEEE 802.3ah, *Ethernet in the First Mile*
IEEE 802.3ba, *40 Gb/s and 100 Gb/s Ethernet*
IEEE 802.3i, *Ethernet*
IEEE 802.3u, *Fast Ethernet*
IEEE 802.3x, *Ethernet Flow Control*
IEEE 802.3z, *Gigabit Ethernet*
ITU-T G.8031/Y.1342, *Ethernet Linear Protection Switching*
ITU-T G.8032/Y.1344, *Ethernet Ring Protection Switching*
ITU-T Y.1731, *OAM functions and mechanisms for Ethernet based networks*

Ethernet VPN (EVPN)

draft-ietf-bess-evpn-ac-df-01, *AC-Influenced Designated Forwarder Election for EVPN*
draft-ietf-bess-evpn-pref-df-01, *Preference-based EVPN DF Election*
draft-ietf-bess-evpn-prefix-advertisement-11, *IP Prefix Advertisement in EVPN*

draft-ietf-bess-evpn-proxy-arp-nd-04, *Operational Aspects of Proxy-ARP/ND in EVPN Networks*

draft-ietf-bess-evpn-vpls-seamless-integ-03, *(PBB-)EVPN Seamless Integration with (PBB-)VPLS*

draft-snr-bess-pbb-evpn-isid-cmacflush-01, *PBB-EVPN ISID-based CMAC-Flush*

RFC 7432, *BGP MPLS-Based Ethernet VPN*

RFC 7623, *Provider Backbone Bridging Combined with Ethernet VPN (PBB-EVPN)*

RFC 8214, *Virtual Private Wire Service Support in Ethernet VPN*

RFC 8317, *Ethernet-Tree (E-Tree) Support in Ethernet VPN (EVPN) an Provider Backbone Bridging EVPN (PBB-EVPN)*

RFC 8365, *A Network Virtualization Overlay Solution Using Ethernet VPN (EVPN)*

Frame Relay

ANSI T1.617 Annex D, *DSS1 - Signalling Specification For Frame Relay Bearer Service*

FRF.1.2, *PVC User-to-Network Interface (UNI) Implementation Agreement*

FRF.12, *Frame Relay Fragmentation Implementation Agreement*

FRF.16.1, *Multilink Frame Relay UNI/NNI Implementation Agreement*

FRF.5, *Frame Relay/ATM PVC Network Interworking Implementation*

FRF2.2, *PVC Network-to-Network Interface (NNI) Implementation Agreement*

ITU-T Q.933 Annex A, *Additional procedures for Permanent Virtual Connection (PVC) status management*

Generalized Multiprotocol Label Switching (GMPLS)

draft-ietf-ccamp-rsvp-te-srlg-collect-04, *RSVP-TE Extensions for Collecting SRLG Information*

RFC 3471, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description*

RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions*

RFC 4204, *Link Management Protocol (LMP)*

RFC 4208, *Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model*

RFC 4872, *RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery*

RFC 5063, *Extensions to GMPLS Resource Reservation Protocol (RSVP) Graceful Restart* (helper mode)

gRPC Remote Procedure Calls (gRPC)

gnmi.proto, *gRPC Network Management Interface (gNMI), version 0.4.0*

gRPC Network Management Interface (gNMI), *Capabilities, Get, Set, Subscribe (ONCE, SAMPLE, ON_CHANGE)*

Intermediate System to Intermediate System (IS-IS)

draft-ietf-isis-mi-02, *IS-IS Multi-Instance*

draft-kaplan-isis-ext-eth-02, *Extended Ethernet Frame Size Support*

ISO/IEC 10589:2002, Second Edition, Nov. 2002, *Intermediate system to Intermediate system intra-domain routeing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)*

RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*

RFC 2973, *IS-IS Mesh Groups*

RFC 3359, *Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System*

RFC 3719, *Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)*

RFC 3787, *Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS)*

RFC 4971, *Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information*

RFC 5120, *M-ISIS: Multi Topology (MT) Routing in IS-IS*

RFC 5130, *A Policy Control Mechanism in IS-IS Using Administrative Tags*

RFC 5301, *Dynamic Hostname Exchange Mechanism for IS-IS*

RFC 5302, *Domain-wide Prefix Distribution with Two-Level IS-IS*

RFC 5303, *Three-Way Handshake for IS-IS Point-to-Point Adjacencies*

RFC 5304, *IS-IS Cryptographic Authentication*

RFC 5305, *IS-IS Extensions for Traffic Engineering TE*

RFC 5306, *Restart Signaling for IS-IS* (helper mode)

RFC 5307, *IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*

RFC 5308, *Routing IPv6 with IS-IS*

RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*

RFC 5310, *IS-IS Generic Cryptographic Authentication*

RFC 6213, *IS-IS BFD-Enabled TLV*

RFC 6232, *Purge Originator Identification TLV for IS-IS*

RFC 6233, *IS-IS Registry Extension for Purges*

RFC 6329, *IS-IS Extensions Supporting IEEE 802.1aq Shortest Path Bridging*

RFC 7775, *IS-IS Route Preference for Extended IP and IPv6 Reachability*

RFC 7794, *IS-IS Prefix Attributes for Extended IPv4 and IPv6 Reachability*

RFC 8202, *IS-IS Multi-Instance* (single topology)

Internet Protocol (IP) — Fast Reroute

draft-ietf-rtgwg-lfa-manageability-08, *Operational management of Loop Free Alternates*

RFC 5286, *Basic Specification for IP Fast Reroute: Loop-Free Alternates*

RFC 7431, *Multicast-Only Fast Reroute*

RFC 7490, *Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)*

Internet Protocol (IP) — General

draft-grant-tacacs-02, *The TACACS+ Protocol*

RFC 768, *User Datagram Protocol*

RFC 793, *Transmission Control Protocol*

RFC 854, *Telnet Protocol Specifications*

RFC 1350, *The TFTP Protocol (revision 2)*

RFC 2347, *TFTP Option Extension*

RFC 2348, *TFTP Blocksize Option*

RFC 2349, *TFTP Timeout Interval and Transfer Size Options*

RFC 2428, *FTP Extensions for IPv6 and NATs*

RFC 2784, *Generic Routing Encapsulation (GRE)*

RFC 2890, *Key and Sequence Number Extensions to GRE*

RFC 4250, *The Secure Shell (SSH) Protocol Assigned Numbers*

RFC 4251, *The Secure Shell (SSH) Protocol Architecture*

RFC 4252, *The Secure Shell (SSH) Authentication Protocol* (publickey, password)

RFC 4253, *The Secure Shell (SSH) Transport Layer Protocol*

RFC 4254, *The Secure Shell (SSH) Connection Protocol*

RFC 4632, *Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan*

RFC 5082, *The Generalized TTL Security Mechanism (GTSM)*

RFC 5656, *Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer* (ECDSA)

RFC 5925, *The TCP Authentication Option*

RFC 5926, *Cryptographic Algorithms for the TCP Authentication Option (TCP-AO)*

RFC 6398, *IP Router Alert Considerations and Usage* (MLD)

RFC 6528, *Defending against Sequence Number Attacks*

Internet Protocol (IP) — Multicast

cisco-ipmulticast/pim-autorp-spec01, *Auto-RP: Automatic discovery of Group-to-RP mappings for IP multicast* (version 1)

draft-dolganow-bess-mvpn-expl-track-01, *Explicit Tracking with Wild Card Routes in Multicast VPN*

draft-ietf-bier-mvpn-11, *Multicast VPN Using BIER*

draft-ietf-idmr-traceroute-ipm-07, *A "traceroute" facility for IP Multicast*

draft-ietf-l2vpn-vpls-pim-snooping-07, *Protocol Independent Multicast (PIM) over Virtual Private LAN Service (VPLS)*

draft-ietf-mboned-mtrace-v2-17, *Mtrace Version 2: Traceroute Facility for IP Multicast*

RFC 1112, *Host Extensions for IP Multicasting*

RFC 2236, *Internet Group Management Protocol, Version 2*

RFC 2365, *Administratively Scoped IP Multicast*

RFC 2375, *IPv6 Multicast Address Assignments*

RFC 2710, *Multicast Listener Discovery (MLD) for IPv6*

RFC 3306, *Unicast-Prefix-based IPv6 Multicast Addresses*

RFC 3376, *Internet Group Management Protocol, Version 3*

RFC 3446, *Anycast Rendevous Point (RP) mechanism using Protocol Independent Multicast (PIM) and Multicast Source Discovery Protocol (MSDP)*

RFC 3590, *Source Address Selection for the Multicast Listener Discovery (MLD) Protocol*

RFC 3618, *Multicast Source Discovery Protocol (MSDP)*

RFC 3810, *Multicast Listener Discovery Version 2 (MLDv2) for IPv6*

RFC 3956, *Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address*

RFC 3973, *Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)* (auto-RP groups)

RFC 4541, *Considerations for Internet Group Management Protocol (IGMP) and Multicast Listener Discovery (MLD) Snooping Switches*

RFC 4604, *Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast*

RFC 4607, *Source-Specific Multicast for IP*

RFC 4608, *Source-Specific Protocol Independent Multicast in 232/8*

RFC 4610, *Anycast-RP Using Protocol Independent Multicast (PIM)*

RFC 4611, *Multicast Source Discovery Protocol (MSDP) Deployment Scenarios*

RFC 5059, *Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)*

RFC 5186, *Internet Group Management Protocol Version 3 (IGMPv3) / Multicast Listener Discovery Version 2 (MLDv2) and Multicast Routing Protocol Interaction*

RFC 5384, *The Protocol Independent Multicast (PIM) Join Attribute Format*

RFC 5496, *The Reverse Path Forwarding (RPF) Vector TLV*

RFC 6037, *Cisco Systems' Solution for Multicast in MPLS/BGP IP VPNs*

RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root*

RFC 6513, *Multicast in MPLS/BGP IP VPNs*

RFC 6514, *BGP Encodings and Procedures for Multicast in MPLS/IP VPNs*

RFC 6515, *IPv4 and IPv6 Infrastructure Addresses in BGP Updates for Multicast VPNs*

RFC 6516, *IPv6 Multicast VPN (MVPN) Support Using PIM Control Plane and Selective Provider Multicast Service Interface (S-PMSI) Join Messages*

RFC 6625, *Wildcards in Multicast VPN Auto-Discover Routes*

RFC 6826, *Multipoint LDP In-Band Signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Path*

RFC 7246, *Multipoint Label Distribution Protocol In-Band Signaling in a Virtual Routing and Forwarding (VRF) Table Context*

RFC 7385, *IANA Registry for P-Multicast Service Interface (PMSI) Tunnel Type Code Points*

RFC 7716, *Global Table Multicast with BGP Multicast VPN (BGP-MVPN) Procedures*

RFC 7761, *Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)*

RFC 8279, *Multicast Using Bit Index Explicit Replication (BIER)*

RFC 8296, *Encapsulation for Bit Index Explicit Replication (BIER) in MPLS and Non-MPLS Networks* (MPLS encapsulation)

RFC 8401, *Bit Index Explicit Replication (BIER) Support via IS-IS*

Internet Protocol (IP) — Version 4

RFC 791, *Internet Protocol*

RFC 792, *Internet Control Message Protocol*

RFC 826, *An Ethernet Address Resolution Protocol*

RFC 951, *Bootstrap Protocol (BOOTP)*

RFC 1034, *Domain Names - Concepts and Facilities*

RFC 1035, *Domain Names - Implementation and Specification*

RFC 1191, *Path MTU Discovery* (router specification)

RFC 1519, *Classless Inter-Domain Routing (CIDR): an Address Assignment and Aggregation Strategy*

RFC 1534, *Interoperation between DHCP and BOOTP*

RFC 1542, *Clarifications and Extensions for the Bootstrap Protocol*

RFC 1812, *Requirements for IPv4 Routers*

RFC 1918, *Address Allocation for Private Internets*

RFC 2003, *IP Encapsulation within IP*

RFC 2131, *Dynamic Host Configuration Protocol*

RFC 2132, *DHCP Options and BOOTP Vendor Extensions*

RFC 2401, *Security Architecture for Internet Protocol*

RFC 3021, *Using 31-Bit Prefixes on IPv4 Point-to-Point Links*

RFC 3046, *DHCP Relay Agent Information Option (Option 82)*

RFC 3768, *Virtual Router Redundancy Protocol (VRRP)*

RFC 4884, *Extended ICMP to Support Multi-Part Messages* (ICMPv4 and ICMPv6 Time Exceeded)

Internet Protocol (IP) — Version 6

RFC 2464, *Transmission of IPv6 Packets over Ethernet Networks*

RFC 2529, *Transmission of IPv6 over IPv4 Domains without Explicit Tunnels*

RFC 3122, *Extensions to IPv6 Neighbor Discovery for Inverse Discovery Specification*

RFC 3315, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*

RFC 3587, *IPv6 Global Unicast Address Format*

RFC 3596, *DNS Extensions to Support IP version 6*

RFC 3633, *IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6*

RFC 3646, *DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)*

RFC 3736, *Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6*

RFC 3971, *SEcure Neighbor Discovery (SEND)*

RFC 3972, *Cryptographically Generated Addresses (CGA)*

RFC 4007, *IPv6 Scoped Address Architecture*

RFC 4193, *Unique Local IPv6 Unicast Addresses*

RFC 4291, *Internet Protocol Version 6 (IPv6) Addressing Architecture*

RFC 4443, *Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification*

RFC 4861, *Neighbor Discovery for IP version 6 (IPv6)*

RFC 4862, *IPv6 Stateless Address Autoconfiguration* (router functions)

RFC 4890, *Recommendations for Filtering ICMPv6 Messages in Firewalls*

RFC 4941, *Privacy Extensions for Stateless Address Autoconfiguration in IPv6*

RFC 5007, *DHCPv6 Leasequery*

RFC 5095, *Deprecation of Type 0 Routing Headers in IPv6*

RFC 5722, *Handling of Overlapping IPv6 Fragments*

RFC 5798, *Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6* (IPv6)

RFC 5952, *A Recommendation for IPv6 Address Text Representation*

RFC 6092, *Recommended Simple Security Capabilities in Customer Premises Equipment (CPE) for Providing Residential IPv6 Internet Service* (Internet Control and Management, Upper-Layer Transport Protocols, UDP Filters, IPsec and Internet Key Exchange (IKE), TCP Filters)

RFC 6106, *IPv6 Router Advertisement Options for DNS Configuration*

RFC 6164, *Using 127-Bit IPv6 Prefixes on Inter-Router Links*

RFC 8021, *Generation of IPv6 Atomic Fragments Considered Harmful*

RFC 8200, *Internet Protocol, Version 6 (IPv6) Specification*

RFC 8201, *Path MTU Discovery for IP version 6*

Internet Protocol Security (IPsec)

draft-ietf-ipsec-isakmp-mode-cfg-05, *The ISAKMP Configuration Method*

draft-ietf-ipsec-isakmp-xauth-06, *Extended Authentication within ISAKMP/Oakley (XAUTH)*

RFC 2401, *Security Architecture for the Internet Protocol*

RFC 2403, *The Use of HMAC-MD5-96 within ESP and AH*

RFC 2404, *The Use of HMAC-SHA-1-96 within ESP and AH*

RFC 2405, *The ESP DES-CBC Cipher Algorithm With Explicit IV*

RFC 2406, *IP Encapsulating Security Payload (ESP)*

RFC 2407, *IPsec Domain of Interpretation for ISAKMP (IPsec DoI)*

RFC 2408, *Internet Security Association and Key Management Protocol (ISAKMP)*

RFC 2409, *The Internet Key Exchange (IKE)*

RFC 2410, *The NULL Encryption Algorithm and Its Use With IPsec*

RFC 3526, *More Modular Exponential (MODP) Diffie-Hellman group for Internet Key Exchange (IKE)*

RFC 3566, *The AES-XCBC-MAC-96 Algorithm and Its Use With IPsec*

RFC 3602, *The AES-CBC Cipher Algorithm and Its Use with IPsec*

RFC 3706, *A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers*

RFC 3947, *Negotiation of NAT-Traversal in the IKE*

RFC 3948, *UDP Encapsulation of IPsec ESP Packets*

RFC 4106, *The Use of Galois/Counter Mode (GCM) in IPsec ESP*

RFC 4210, *Internet X.509 Public Key Infrastructure Certificate Management Protocol (CMP)*

RFC 4211, *Internet X.509 Public Key Infrastructure Certificate Request Message Format (CRMF)*

RFC 4301, *Security Architecture for the Internet Protocol*

RFC 4303, *IP Encapsulating Security Payload*

RFC 4307, *Cryptographic Algorithms for Use in the Internet Key Exchange Version 2 (IKEv2)*

RFC 4308, *Cryptographic Suites for IPsec*

RFC 4434, *The AES-XCBC-PRF-128 Algorithm for the Internet Key Exchange Protocol (IKE)*

RFC 4543, *The Use of Galois Message Authentication Code (GMAC) in IPsec ESP and AH*

RFC 4868, *Using HMAC-SHA-256, HMAC-SHA-384, and HMAC-SHA-512 with IPSec*

RFC 4945, *The Internet IP Security PKI Profile of IKEv1/ISAKMP, IKEv2 and PKIX*

RFC 5019, *The Lightweight Online Certificate Status Protocol (OCSP) Profile for High-Volume Environments*

RFC 5280, *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile*

RFC 5282, *Using Authenticated Encryption Algorithms with the Encrypted Payload of the IKEv2 Protocol*

RFC 5903, *ECP Groups for IKE and IKEv2*

RFC 5998, *An Extension for EAP-Only Authentication in IKEv2*

RFC 6379, *Suite B Cryptographic Suites for IPsec*

RFC 6380, *Suite B Profile for Internet Protocol Security (IPsec)*

RFC 6712, *Internet X.509 Public Key Infrastructure -- HTTP Transfer for the Certificate Management Protocol (CMP)*

RFC 6960, *X.509 Internet Public Key Infrastructure Online Certificate Status Protocol - OCSP*

RFC 7296, *Internet Key Exchange Protocol Version 2 (IKEv2)*

RFC 7321, *Cryptographic Algorithm Implementation Requirements and Usage Guidance for Encapsulating Security Payload (ESP) and Authentication Header (AH)*

RFC 7383, *Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation*

RFC 7427, *Signature Authentication in the Internet Key Exchange Version 2 (IKEv2)*

RFC 7468, *Textual Encodings of PKIX, PKCS, and CMS Structures*

Label Distribution Protocol (LDP)

draft-ietf-mpls-ldp-ip-pw-capability-09, *Controlling State Advertisements Of Non-negotiated LDP Applications*

draft-pdutta-mpls-ldp-adj-capability-00, *LDP Adjacency Capabilities*

draft-pdutta-mpls-ldp-v2-00, *LDP Version 2*

draft-pdutta-mpls-mldp-up-redundancy-00, *Upstream LSR Redundancy for Multi-point LDP Tunnels*

draft-pdutta-mpls-multi-ldp-instance-00, *Multiple LDP Instances*

draft-pdutta-mpls-tldp-hello-reduce-04, *Targeted LDP Hello Reduction*

RFC 3037, *LDP Applicability*

RFC 3478, *Graceful Restart Mechanism for Label Distribution Protocol* (helper mode)

RFC 5036, *LDP Specification*

RFC 5283, *LDP Extension for Inter-Area Label Switched Paths (LSPs)*

RFC 5443, *LDP IGP Synchronization*

RFC 5561, *LDP Capabilities*

RFC 5919, *Signaling LDP Label Advertisement Completion*

RFC 6388, *Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*

RFC 6512, *Using Multipoint LDP When the Backbone Has No Route to the Root*

RFC 6826, *Multipoint LDP in-band signaling for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths*

RFC 7032, *LDP Downstream-on-Demand in Seamless MPLS*

RFC 7552, *Updates to LDP for IPv6*

Layer Two Tunneling Protocol (L2TP) Network Server (LNS)

draft-mammoliti-l2tp-accessline-avp-04, *Layer 2 Tunneling Protocol (L2TP) Access Line Information Attribute Value Pair (AVP) Extensions*

RFC 2661, *Layer Two Tunneling Protocol "L2TP"*

RFC 2809, *Implementation of L2TP Compulsory Tunneling via RADIUS*

RFC 3438, *Layer Two Tunneling Protocol (L2TP) Internet Assigned Numbers: Internet Assigned Numbers Authority (IANA) Considerations Update*

RFC 3931, *Layer Two Tunneling Protocol - Version 3 (L2TPv3)*

RFC 4719, *Transport of Ethernet Frames over Layer 2 Tunneling Protocol Version 3 (L2TPv3)*

RFC 4951, *Fail Over Extensions for Layer 2 Tunneling Protocol (L2TP) "failover"*

Management

draft-ieft-snmpv3-update-mib-05, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*

draft-ietf-isis-wg-mib-06, *Management Information Base for Intermediate System to Intermediate System (IS-IS)*

draft-ietf-mboned-msdp-mib-01, *Multicast Source Discovery protocol MIB*

draft-ietf-mpls-ldp-mib-07, *Definitions of Managed Objects for the Multiprotocol Label Switching, Label Distribution Protocol (LDP)*

draft-ietf-mpls-lsr-mib-06, *Multiprotocol Label Switching (MPLS) Label Switching Router (LSR) Management Information Base Using SMIv2*

draft-ietf-mpls-te-mib-04, *Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base*

draft-ietf-ospf-mib-update-08, *OSPF Version 2 Management Information Base*

draft-ietf-vrrp-unified-mib-06, *Definitions of Managed Objects for the VRRP over IPv4 and IPv6* (IPv6)

ianaaddressfamilynumbers-mib, *IANA-ADDRESS-FAMILY-NUMBERS-MIB*

ianagmplstc-mib, *IANA-GMPLS-TC-MIB*

ianaiftype-mib, *IANAifType-MIB*

ianaiprouteprotocol-mib, *IANA-RTPROTO-MIB*

IEEE8021-CFM-MIB, *IEEE P802.1ag(TM) CFM MIB*

IEEE8021-PAE-MIB, *IEEE 802.1X MIB*

IEEE8023-LAG-MIB, *IEEE 802.3ad MIB*

LLDP-MIB, *IEEE P802.1AB(TM) LLDP MIB*

openconfig-bgp.yang version 3.0.1, *BGP Module*

openconfig-bgp-common.yang version 3.0.1, *BGP Common Module*

openconfig-bgp-common-multiprotocol.yang version 3.0.1, *BGP Common Multiprotocol Module*

openconfig-bgp-common-structure.yang version 3.0.1, *BGP Common Structure Module*

openconfig-bgp-global.yang version 3.0.1, *BGP Global Module*

openconfig-bgp-neighbor.yang version 3.0.1, *BGP Neighbor Module*

openconfig-bgp-peer-group.yang version 3.0.1, *BGP Peer Group Module*

openconfig-bgp-policy.yang version 4.0.1, *BGP Policy Module*

openconfig-if-aggregate.yang version 2.0.0, *Interfaces Aggregated Model*

openconfig-if-ethernet.yang version 2.0.0, *Interfaces Ethernet Model*

openconfig-if-ip.yang version 2.0.0, *Interfaces IP Module*

openconfig-if-ip-ext.yang version 2.0.0, *Interfaces IP Extensions Module*

openconfig-interfaces.yang version 2.0.0, *Interfaces Module*

openconfig-isis.yang version 0.3.0, *IS-IS Module*

openconfig-isis-lsp.yang version 0.3.0, *IS-IS LSP Module*

openconfig-isis-routing.yang version 0.3.0, *IS-IS Routing Module*

openconfig-lacp.yang version 1.1.0, *LACP Module*

openconfig-lldp.yang version 0.1.0, *LLDP Module*

openconfig-local-routing.yang version 1.0.1, *Local Routing Module*

openconfig-network-instance.yang version 0.8.0, *Network Instance Module*

openconfig-routing-policy.yang version 3.0.0, *Routing Policy Module*

openconfig-vlan.yang version 2.0.0, *VLAN Module*

RFC 1157, *A Simple Network Management Protocol (SNMP)*

RFC 1212, *Concise MIB Definitions*

RFC 1213, *Management Information Base for Network Management of TCP/IP-based Internets: MIB-II*

RFC 1215, *A Convention for Defining Traps for use with the SNMP*

RFC 1724, *RIP Version 2 MIB Extension*

RFC 1901, *Introduction to Community-based SNMPv2*

RFC 2021, *Remote Network Monitoring Management Information Base Version 2 using SMIv2*

RFC 2115, *Management Information Base for Frame Relay DTEs Using SMIv2*

RFC 2206, *RSVP Management Information Base using SMIv2*

RFC 2213, *Integrated Services Management Information Base using SMIv2*

RFC 2494, *Definitions of Managed Objects for the DS0 and DS0 Bundle Interface Type*

RFC 2514, *Definitions of Textual Conventions and OBJECT-IDENTITIES for ATM Management*

RFC 2515, *Definitions of Managed Objects for ATM Management*

RFC 2570, *SNMP Version 3 Framework*

RFC 2571, *An Architecture for Describing SNMP Management Frameworks*

RFC 2572, *Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)*

RFC 2573, *SNMP Applications*

RFC 2574, *User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)*

RFC 2575, *View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)*

RFC 2578, *Structure of Management Information Version 2 (SMIv2)*

RFC 2579, *Textual Conventions for SMIv2*

RFC 2580, *Conformance Statements for SMIv2*

RFC 2787, *Definitions of Managed Objects for the Virtual Router Redundancy Protocol*

RFC 2819, *Remote Network Monitoring Management Information Base*

RFC 2856, *Textual Conventions for Additional High Capacity Data Types*

RFC 2863, *The Interfaces Group MIB*

RFC 2864, *The Inverted Stack Table Extension to the Interfaces Group MIB*

RFC 2933, *Internet Group Management Protocol MIB*

RFC 3014, *Notification Log MIB*

RFC 3164, *The BSD syslog Protocol*

RFC 3165, *Definitions of Managed Objects for the Delegation of Management Scripts*

RFC 3231, *Definitions of Managed Objects for Scheduling Management Operations*

RFC 3273, *Remote Network Monitoring Management Information Base for High Capacity Networks*

RFC 3416. *Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)*

RFC 3417, *Transport Mappings for the Simple Network Management Protocol (SNMP)* (SNMP over UDP over IPv4)

RFC 3419, *Textual Conventions for Transport Addresses*

RFC 3498, *Definitions of Managed Objects for Synchronous Optical Network (SONET) Linear Automatic Protection Switching (APS) Architectures*

RFC 3584, *Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network Management Framework*

RFC 3592, *Definitions of Managed Objects for the Synchronous Optical Network/ Synchronous Digital Hierarchy (SONET/SDH) Interface Type*

RFC 3593, *Textual Conventions for MIB Modules Using Performance History Based on 15 Minute Intervals*

RFC 3635, *Definitions of Managed Objects for the Ethernet-like Interface Types*

RFC 3637, *Definitions of Managed Objects for the Ethernet WAN Interface Sublayer*

RFC 3826, *The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model*

RFC 3877, *Alarm Management Information Base (MIB)*

RFC 3895, *Definitions of Managed Objects for the DS1, E1, DS2, and E2 Interface Types*

RFC 3896, *Definitions of Managed Objects for the DS3/E3 Interface Type*

RFC 4001, *Textual Conventions for Internet Network Addresses*

RFC 4022, *Management Information Base for the Transmission Control Protocol (TCP)*

RFC 4113, *Management Information Base for the User Datagram Protocol (UDP)*

RFC 4220, *Traffic Engineering Link Management Information Base*

RFC 4273, *Definitions of Managed Objects for BGP-4*

RFC 4292, *IP Forwarding Table MIB*

RFC 4293, *Management Information Base for the Internet Protocol (IP)*

RFC 4379, *Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*

RFC 4511, *Lightweight Directory Access Protocol (LDAP): The Protocol*

RFC 4513, *Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms* (TLS)

RFC 4631, *Link Management Protocol (LMP) Management Information Base (MIB)*

RFC 4878, *Definitions and Managed Objects for Operations, Administration, and Maintenance (OAM) Functions on Ethernet-Like Interfaces*

RFC 5101, *Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information*

RFC 5102, *Information Model for IP Flow Information Export*

RFC 5246, *The Transport Layer Security (TLS) Protocol Version 1.2* (TLS client, RSA public key)

RFC 6424, *Mechanism for Performing Label Switched Path Ping (LSP Ping) over MPLS Tunnels*

RFC 6425, *Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping*

RFC 6991, *Common YANG Data Types*

RFC 7420, *Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module*

RFC 7950, *The YANG 1.1 Data Modeling Language*

SFLOW-MIB, *sFlow MIB Version 1.3 (Draft 5)*

Multiprotocol Label Switching — Transport Profile (MPLS-TP)

RFC 5586, *MPLS Generic Associated Channel*

RFC 5921, *A Framework for MPLS in Transport Networks*

RFC 5960, *MPLS Transport Profile Data Plane Architecture*

RFC 6370, *MPLS Transport Profile (MPLS-TP) Identifiers*

RFC 6378, *MPLS Transport Profile (MPLS-TP) Linear Protection*

RFC 6426, *MPLS On-Demand Connectivity and Route Tracing*

RFC 6427, *MPLS Fault Management Operations, Administration, and Maintenance (OAM)*

RFC 6428, *Proactive Connectivity Verification, Continuity Check and Remote Defect indication for MPLS Transport Profile*

RFC 6478, *Pseudowire Status for Static Pseudowires*

RFC 7213, *MPLS Transport Profile (MPLS-TP) Next-Hop Ethernet Addressing*

Multiprotocol Label Switching (MPLS)

draft-ietf-teas-sr-rsvp-coexistence-rec-02, *Recommendations for RSVP-TE and Segment Routing LSP co-existence*

RFC 3031, *Multiprotocol Label Switching Architecture*

RFC 3032, *MPLS Label Stack Encoding*

RFC 3443, *Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks*

RFC 4023, *Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)*

RFC 4182, *Removing a Restriction on the use of MPLS Explicit NULL*

RFC 5332, *MPLS Multicast Encapsulations*

RFC 5884, *Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)*

RFC 6790, *The Use of Entropy Labels in MPLS Forwarding*

RFC 7510, *Encapsulating MPLS in UDP*

Network Address Translation (NAT)

draft-ietf-behave-address-format-10, *IPv6 Addressing of IPv4/IPv6 Translators*

draft-ietf-behave-v6v4-xlate-23, *IP/ICMP Translation Algorithm*

draft-miles-behave-l2nat-00, *Layer2-Aware NAT*

draft-nishitani-cgn-02, *Common Functions of Large Scale NAT (LSN)*

RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP*

RFC 5382, *NAT Behavioral Requirements for TCP*

RFC 5508, *NAT Behavioral Requirements for ICMP*

RFC 6146, *Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers*

RFC 6333, *Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion*

RFC 6334, *Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite*

RFC 6887, *Port Control Protocol (PCP)*

RFC 6888, *Common Requirements For Carrier-Grade NATs (CGNs)*

RFC 7915, *IP/ICMP Translation Algorithm*

Network Configuration Protocol (NETCONF)

RFC 5277, *NETCONF Event Notifications*

RFC 6022, *YANG Module for NETCONF Monitoring*

RFC 6241, *Network Configuration Protocol (NETCONF)*

RFC 6242, *Using the NETCONF Protocol over Secure Shell (SSH)*

RFC 6243, *With-defaults Capability for NETCONF*

RFC 7895, *YANG Module Library*

Open Shortest Path First (OSPF)

draft-ietf-ospf-ospfv3-lsa-extend-13, *OSPFv3 LSA Extendibility*

RFC 1586, *Guidelines for Running OSPF Over Frame Relay Networks*

RFC 1765, *OSPF Database Overflow*

RFC 2328, *OSPF Version 2*

RFC 3101, *The OSPF Not-So-Stubby Area (NSSA) Option*

RFC 3509, *Alternative Implementations of OSPF Area Border Routers*

RFC 3623, *Graceful OSPF Restart Graceful OSPF Restart* (helper mode)

RFC 3630, *Traffic Engineering (TE) Extensions to OSPF Version 2*

RFC 4203, *OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)*

RFC 4222, *Prioritized Treatment of Specific OSPF Version 2 Packets and Congestion Avoidance*

RFC 4552, *Authentication/Confidentiality for OSPFv3*

RFC 4576, *Using a Link State Advertisement (LSA) Options Bit to Prevent Looping in BGP/MPLS IP Virtual Private Networks (VPNs)*

RFC 4577, *OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)*

RFC 5185, *OSPF Multi-Area Adjacency*

RFC 5187, *OSPFv3 Graceful Restart* (helper mode)

RFC 5243, *OSPF Database Exchange Summary List Optimization*

RFC 5250, *The OSPF Opaque LSA Option*

RFC 5309, *Point-to-Point Operation over LAN in Link State Routing Protocols*

RFC 5340, *OSPF for IPv6*

RFC 5709, *OSPFv2 HMAC-SHA Cryptographic Authentication*

RFC 5838, *Support of Address Families in OSPFv3*

RFC 6987, *OSPF Stub Router Advertisement*

RFC 7684, *OSPFv2 Prefix/Link Attribute Advertisement*

RFC 7770, *Extensions to OSPF for Advertising Optional Router Capabilities*

OpenFlow

TS-007, *OpenFlow Switch Specification Version 1.3.1* (OpenFlow-hybrid switches)

Path Computation Element Protocol (PCEP)

draft-alvarez-pce-path-profiles-04, *PCE Path Profiles*

draft-ietf-pce-segment-routing-08, *PCEP Extensions for Segment Routing*

draft-ietf-pce-stateful-pce-14, *PCEP Extensions for Stateful PCE*

RFC 5440, *Path Computation Element (PCE) Communication Protocol (PCEP)*

Point-to-Point Protocol (PPP)

RFC 1332, *The PPP Internet Protocol Control Protocol (IPCP)*

RFC 1377, *The PPP OSI Network Layer Control Protocol (OSINLCP)*

RFC 1661, *The Point-to-Point Protocol (PPP)*

RFC 1662, *PPP in HDLC-like Framing*

RFC 1877, *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*

RFC 1989, *PPP Link Quality Monitoring*

RFC 1990, *The PPP Multilink Protocol (MP)*

RFC 1994, *PPP Challenge Handshake Authentication Protocol (CHAP)*

RFC 2153, *PPP Vendor Extensions*

RFC 2516, *A Method for Transmitting PPP Over Ethernet (PPPoE)*

RFC 2615, *PPP over SONET/SDH*

RFC 2686, *The Multi-Class Extension to Multi-Link PPP*

RFC 2878, *PPP Bridging Control Protocol (BCP)*

RFC 4638, *Accommodating a Maximum Transit Unit/Maximum Receive Unit (MTU/MRU) Greater Than 1492 in the Point-to-Point Protocol over Ethernet (PPPoE)*

RFC 5072, *IP Version 6 over PPP*

Policy Management and Credit Control

3GPP TS 29.212 Release 11, *Policy and Charging Control (PCC); Reference points* (Gx support as it applies to wireline environment (BNG))

RFC 3588, *Diameter Base Protocol*

RFC 4006, *Diameter Credit-Control Application*

Pseudowire

draft-ietf-l2vpn-vpws-iw-oam-04, *OAM Procedures for VPWS Interworking*

MFA Forum 12.0.0, *Multiservice Interworking - Ethernet over MPLS*

MFA Forum 13.0.0, *Fault Management for Multiservice Interworking v1.0*

MFA Forum 16.0.0, *Multiservice Interworking - IP over MPLS*

MFA Forum 9.0.0, *The Use of Virtual trunks for ATM/MPLS Control Plane Interworking*

RFC 3916, *Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)*

RFC 3985, *Pseudo Wire Emulation Edge-to-Edge (PWE3)*

RFC 4385, *Pseudo Wire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN*

RFC 4446, *IANA Allocations for Pseudowire Edge to Edge Emulation (PWE3)*

RFC 4447, *Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)*

RFC 4448, *Encapsulation Methods for Transport of Ethernet over MPLS Networks*

RFC 4619, *Encapsulation Methods for Transport of Frame Relay over Multiprotocol Label Switching (MPLS) Networks*

RFC 4717, *Encapsulation Methods for Transport Asynchronous Transfer Mode (ATM) over MPLS Networks*

RFC 4816, *Pseudowire Emulation Edge-to-Edge (PWE3) Asynchronous Transfer Mode (ATM) Transparent Cell Transport Service*

RFC 5085, *Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires*

RFC 5659, *An Architecture for Multi-Segment Pseudowire Emulation Edge-to-Edge*

RFC 5885, *Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)*

RFC 6073, *Segmented Pseudowire*

RFC 6310, *Pseudowire (PW) Operations, Administration, and Maintenance (OAM) Message Mapping*

RFC 6391, *Flow-Aware Transport of Pseudowires over an MPLS Packet Switched Network*

RFC 6575, *Address Resolution Protocol (ARP) Mediation for IP Interworking of Layer 2 VPNs*

RFC 6718, *Pseudowire Redundancy*

RFC 6829, *Label Switched Path (LSP) Ping for Pseudowire Forwarding Equivalence Classes (FECs) Advertised over IPv6*

RFC 6870, *Pseudowire Preferential Forwarding Status bit*

RFC 7023, *MPLS and Ethernet Operations, Administration, and Maintenance (OAM) Interworking*

RFC 7267, *Dynamic Placement of Multi-Segment Pseudowires*

Quality of Service (QoS)

RFC 2430, *A Provider Architecture for Differentiated Services and Traffic Engineering (PASTE)*

RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*

RFC 2598, *An Expedited Forwarding PHB*

RFC 3140, *Per Hop Behavior Identification Codes*

RFC 3260, *New Terminology and Clarifications for Diffserv*

Remote Authentication Dial In User Service (RADIUS)

RFC 2865, *Remote Authentication Dial In User Service (RADIUS)*

RFC 2866, *RADIUS Accounting*

RFC 2867, *RADIUS Accounting Modifications for Tunnel Protocol Support*

RFC 2868, *RADIUS Attributes for Tunnel Protocol Support*

RFC 2869, *RADIUS Extensions*

RFC 3162, *RADIUS and IPv6*

RFC 4818, *RADIUS Delegated-IPv6-Prefix Attribute*

RFC 5176, *Dynamic Authorization Extensions to RADIUS*

RFC 6911, *RADIUS attributes for IPv6 Access Networks*

RFC 6929, *Remote Authentication Dial-In User Service (RADIUS) Protocol Extensions*

Resource Reservation Protocol — Traffic Engineering (RSVP-TE)

draft-newton-mpls-te-dynamic-overbooking-00, *A Diffserv-TE Implementation Model to dynamically change booking factors during failure events*

RFC 2702, *Requirements for Traffic Engineering over MPLS*

RFC 2747, *RSVP Cryptographic Authentication*

RFC 2961, *RSVP Refresh Overhead Reduction Extensions*

RFC 3097, *RSVP Cryptographic Authentication -- Updated Message Type Value*

RFC 3209, *RSVP-TE: Extensions to RSVP for LSP Tunnels*

RFC 3473, *Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions* (IF_ID RSVP_HOP object with unnumbered interfaces and RSVP-TE graceful restart helper procedures)

RFC 3477, *Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)*

RFC 3564, *Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering*

RFC 3906, *Calculating Interior Gateway Protocol (IGP) Routes Over Traffic Engineering Tunnels*

RFC 4090, *Fast Reroute Extensions to RSVP-TE for LSP Tunnels*

RFC 4124, *Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering*

RFC 4125, *Maximum Allocation Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*

RFC 4127, *Russian Dolls Bandwidth Constraints Model for Diffserv-aware MPLS Traffic Engineering*

RFC 4561, *Definition of a Record Route Object (RRO) Node-Id Sub-Object*

RFC 4875, *Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)*

RFC 4950, *ICMP Extensions for Multiprotocol Label Switching*

RFC 5151, *Inter-Domain MPLS and GMPLS Traffic Engineering -- Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Extensions*

RFC 5712, *MPLS Traffic Engineering Soft Preemption*

RFC 5817, *Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks*

Routing Information Protocol (RIP)

RFC 1058, *Routing Information Protocol*

RFC 2080, *RIPng for IPv6*

RFC 2082, *RIP-2 MD5 Authentication*

RFC 2453, *RIP Version 2*

Segment Routing (SR)

draft-filsfils-spring-segment-routing-policy-05, *Segment Routing Policy for Traffic Engineering*

draft-francois-rtgwg-segment-routing-ti-lfa-04, *Topology Independent Fast Reroute using Segment Routing*

draft-gredler-idr-bgp-ls-segment-routing-ext-03, *BGP Link-State extensions for Segment Routing*

draft-ietf-idr-segment-routing-te-policy-02, *Advertising Segment Routing Policies in BGP*

draft-ietf-isis-segment-routing-extensions-04, *IS-IS Extensions for Segment Routing*

draft-ietf-mpls-spring-lsp-ping-02, *Label Switched Path (LSP) Ping/Trace for Segment Routing Networks Using MPLS Dataplane*

draft-ietf-ospf-segment-routing-extensions-04, *OSPF Extensions for Segment Routing*

draft-ietf-spring-conflict-resolution-05, *Segment Routing MPLS Conflict Resolution*

draft-ietf-spring-segment-routing-ldp-interop-09, *Segment Routing interworking with LDP*

Synchronous Optical Networking (SONET)/Synchronous Digital Hierarchy (SDH)

ANSI T1.105.03, *Jitter Network Interfaces*

ANSI T1.105.06, *Physical Layer Specifications*

ANSI T1.105.09, *Network Timing and Synchronization*

ITU-T G.703, *Physical/electrical characteristics of hierarchical digital interfaces*

ITU-T G.707, *Network node interface for the synchronous digital hierarchy (SDH)*

ITU-T G.813, *Timing characteristics of SDH equipment slave clocks (SEC)*

ITU-T G.823, *The control of jitter and wander within digital networks which are based on the 2048 kbit/s hierarchy*

ITU-T G.824, *The control of jitter and wander within digital networks which are based on the 1544 kbit/s hierarchy*

ITU-T G.825, *The control of jitter and wander within digital networks which are based on the synchronous digital hierarchy (SDH)*

ITU-T G.841, *Types and Characteristics of SDH Networks Protection Architecture, issued in October 1998 and as augmented by Corrigendum 1, issued in July 2002*

ITU-T G.957, *Optical interfaces for equipments and systems relating to the synchronous digital hierarchy*

Time Division Multiplexing (TDM)

ANSI T1.403, *DS1 Metallic Interface Specification*

ANSI T1.404, *DS3 Metallic Interface Specification*

Timing

GR-1244-CORE, *Clocks for the Synchronized Network: Common Generic Criteria, Issue 3, May 2005*

GR-253-CORE, *SONET Transport Systems: Common Generic Criteria. Issue 3, September 2000*

IEEE 1588-2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*

ITU-T G.781, *Synchronization layer functions, issued 09/2008*

ITU-T G.813, *Timing characteristics of SDH equipment slave clocks (SEC), issued 03/2003*

ITU-T G.8261, *Timing and synchronization aspects in packet networks, issued 04/2008*

ITU-T G.8262, *Timing characteristics of synchronous Ethernet equipment slave clock (EEC), issued 08/2007*

ITU-T G.8264, *Distribution of timing information through packet networks, issued 10/2008*

ITU-T G.8265.1, *Precision time protocol telecom profile for frequency synchronization, issued 10/2010*

ITU-T G.8275.1, *Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, issued 07/2014*

RFC 5905, *Network Time Protocol Version 4: Protocol and Algorithms Specification*

Two-Way Active Measurement Protocol (TWAMP)

RFC 5357, *A Two-Way Active Measurement Protocol (TWAMP)* (server, unauthenticated mode)

RFC 5938, *Individual Session Control Feature for the Two-Way Active Measurement Protocol (TWAMP)*

RFC 6038, *Two-Way Active Measurement Protocol (TWAMP) Reflect Octets and Symmetrical Size Features*

Virtual Private LAN Service (VPLS)

RFC 4761, *Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling*

RFC 4762, *Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling*

RFC 5501, *Requirements for Multicast Support in Virtual Private LAN Services*

RFC 6074, *Provisioning, Auto-Discovery, and Signaling in Layer 2 Virtual Private Networks (L2VPNs)*

RFC 7041, *Extensions to the Virtual Private LAN Service (VPLS) Provider Edge (PE) Model for Provider Backbone Bridging*

RFC 7117, *Multicast in Virtual Private LAN Service (VPLS)*

Voice and Video

DVB BlueBook A86, *Transport of MPEG-2 TS Based DVB Services over IP Based Networks*

ETSI TS 101 329-5 Annex E, *QoS Measurement for VoIP - Method for determining an Equipment Impairment Factor using Passive Monitoring*

ITU-T G.1020 Appendix I, *Performance Parameter Definitions for Quality of Speech and other Voiceband Applications Utilizing IP Networks - Mean Absolute Packet Delay Variation & Markov Models*

ITU-T G.107, *The E Model - A computational model for use in planning*

ITU-T P.564, *Conformance testing for voice over IP transmission quality assessment models*

RFC 3550 Appendix A.8, *RTP: A Transport Protocol for Real-Time Applications* (estimating the interarrival jitter)

RFC 4585, *Extended RTP Profile for Real-time Transport Control Protocol (RTCP)-Based Feedback (RTP/AVPF)*

RFC 4588, *RTP Retransmission Payload Format*

Wireless Local Area Network (WLAN) Gateway

3GPP TS 23.402, *Architecture enhancements for non-3GPP accesses* (S2a roaming based on GPRS)

Customer Document and Product Support

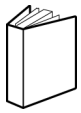

Customer Documentation

Customer Documentation Welcome Page

Technical Support

Product Support Portal

Documentation Feedback

Customer Documentation Feedback

© 2018 Nokia.
3HE 14143 AAAB TQZZA 01